



HD28
.M414

no.
3653-
94

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

IS THE BUSINESS CYCLE A NECESSARY
CONSEQUENCE OF STOCHASTIC GROWTH?

Julio J. Rotemberg
Massachusetts Institute of Technology

Michael Woodford
University of Chicago

WP#3653-94-EFA February 1994

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

IS THE BUSINESS CYCLE A NECESSARY
CONSEQUENCE OF STOCHASTIC GROWTH?

Julio J. Rotemberg
Massachusetts Institute of Technology

Michael Woodford
University of Chicago

WP#3653-94-EFA February 1994

M.I.T. LIBRARIES

FEB 25 1994

RECEIVED

Is the Business Cycle a Necessary Consequence of Stochastic Growth? *

Julio J. Rotemberg
Massachusetts Institute of Technology

Michael Woodford
University of Chicago

February 2, 1994

Abstract

We compute the forecastable changes in output, consumption, and hours implied by a VAR that includes the growth rate of private value added, the share of output that is consumed, and the detrended level of private hours. We show that the size of the forecastable changes in output greatly exceeds that predicted by a standard stochastic growth model, of the kind studied by real business cycle theorists. Contrary to the model's implications, forecastable movements in labor productivity are small and only weakly related to forecasted changes in output. Also, forecasted movements in investment and hours are positively correlated with forecasted movements in output. Finally, and again in contrast to what the growth model implies, forecasted output movements are positively related to the current level of the consumption share and negatively related to the level of hours. We also show that these contrasts between the model and the observations are robust to allowance for measurement error and a variety of other types of transitory disturbances.

*We wish to thank Ludwig Chincarini for research assistance, John Cochrane, John Huizinga, Lucrezia Reichlin, and Argia Sbordonc for helpful discussions, and the NSF for research support.

1 Introduction

The recent literature has given considerable attention to the hypothesis that fluctuations in aggregate economic activity result from stochastic variations in the rate of technical progress (Kydland and Prescott, 1982; Prescott, 1986; King, Plosser and Rebelo, 1988a, 1988b; Plosser, 1989). One of the most appealing features of this “real business cycle” (RBC) hypothesis is its parsimony – it is proposed that the same exogenous changes in the available production technology that determine the long-run changes in output per head also account for short-run variations in output and employment. Rather than a puzzle to be explained through the invocation of a complex mechanism that is introduced into one’s model of the economy solely for that purpose, the business cycle is actually a necessary consequence of stochastic growth. It would be predicted to occur even in the absence of any “frictions” not present in a standard neoclassical model of long-run growth, and in the absence of any other disturbances to the economy, once one recognizes that the technical progress responsible for long run growth is itself stochastic.¹

The demonstration by Nelson and Plosser (1982) that U.S. real GNP has a “unit root” or stochastic trend, rather than exhibiting only transitory fluctuations around a deterministic growth trend, greatly increased the credence given to the real business cycle hypothesis. It is widely accepted that shocks that result in permanent increases in the level of real GNP can only plausibly be interpreted as permanent productivity improvements.² Here we wish to consider whether accepting that there is a stochastic trend in the aggregate production technology requires one to believe also that the business cycle is largely due to (and, in essence, an efficient response to) stochastic productivity growth.³ This issue has been addressed in the RBC literature by computing the predictions of a “calibrated” stochastic growth model for the variability of output, hours worked, and other aggregate quantities, assuming technology shocks of the size indicated by some measure of the variability of technical progress. The question posed is whether even in the absence of other stochastic disturbances to the economy, one would expect to see business cycles of the size and character observed. Typically, such exercises have concluded that technology shocks alone would predict variations in output

¹Thus Prescott (1986) writes: “These [business cycle] observations should not be puzzling, for they are what standard economic theory predicts. For the United States, in fact, ... given the nature of the changing production possibility set, what would be puzzling is if the economy did not display these large fluctuations in output and employment....”

²See, e.g., Blanchard and Fischer (1989, ch. 1).

³We will, for purposes of the inquiry, take at face value the evidence just alluded to. Thus we will not here question whether one can really be certain that U.S. real GNP has a unit root, or whether stochastic productivity growth should be regarded as an exogenous process to which the pattern of resource allocation must adapt, rather than an endogenous consequence of changes in the pattern of resource allocation that may be ultimately caused by shocks of another sort.

of roughly the magnitude observed (see, for example, Plosser 1989). They have also concluded that these fluctuations have the right characteristics, in the sense that the relative volatility of quantities such as aggregate consumption and investment spending, and their degree of co-movement with output fluctuations, correspond roughly to the predictions of the model. Thus it is argued that, granting the existence of the stochastic technical progress and the essential correctness of the neoclassical growth model used to generate the predicted effects of technology shocks on aggregate variables, one can conclude that technology shocks account for a large part of the aggregate fluctuations that are observed. Moreover, it is suggested that such additional variation in aggregate quantities as may be due to other independent disturbances makes little difference for the overall character of business cycles.

Here we undertake a similar exercise, but using a different diagnostic for whether the model can account for the kind of business cycles that we observe. We argue that an important feature of observed business cycles is variation in the rate at which output can be *forecasted* to grow in the future. Indeed, it seems to us that this is an essential aspect of what people mean when they refer to “business cycles” – at different points in time (different “phases of the cycle” in the language popularized by Wesley Mitchell) the outlook for the economy is different. Note that output growth could exhibit substantial variation even in the absence of cycles of this sort. Output might be a random walk, so that the probability distribution over possible future growth patterns would at all times be the same. We show in the next section that this is not true for aggregate fluctuations in the postwar U.S. – instead, there are significant fluctuations in the *forecastable* future growth in output and other aggregate quantities.⁴ Indeed, over a horizon of two to three years, we show that more than half of the variation in output growth is forecastable (even using only a very small set of forecasting variables).

We then ask whether a stochastic growth model would predict fluctuations in the forecastable changes in aggregate quantities of the kind observed. We show that, at least for a popular variant of the model often used in the RBC literature, the answer is absolutely not.⁵ Under a standard calibration of the model

⁴We do not, of course, claim originality for this observation. Since Nelson and Plosser first directed attention to the extent to which aggregate output was similar to a random walk, and argued that this suggested an important role for “supply shocks” in the generation of observed fluctuations, many authors have shown that output is not in fact a random walk (even if it has a “unit root”). The point of the results reported here is to *characterize* the forecastable fluctuations in a way that facilitates comparison with the predictions of a real business cycle model.

⁵Our “baseline” model, including the parameter calibration that we use, is essentially identical to the model with a random walk in technology considered by Plosser (1989) and by King, Plosser and Rebelo (1988b) [hereafter, K-P-R]. The model of Christiano and Eichenbaum (1992) is also quite similar, though certain of their parameter values are rather different. We also discuss below the consequences of variation in the most important of these parameters.

parameters, the variability of the forecastable changes in output predicted by the model is much smaller than that which is indicated by a VAR model of the U.S. data. Essentially, the model predicts that output should be much closer to a random walk than is actually the case.

This finding is related to an observation by Watson (1993), who shows that the RBC model cannot explain the peak of the spectrum of output growth at business cycle frequencies. It is also related to the findings of Cogley and Nason (1993), who criticize several variant RBC models on the grounds that the models cannot account for the observed degree of serial correlation of output growth. We believe that we have identified an even greater discrepancy between the model and the data by analyzing several variables simultaneously. For one thing, as has been stressed by Cochrane and Sbordone (1988), Cochrane (1994a), and Evans and Reichlin (1993), estimates of the forecastability of output growth are greatly increased by the use of a multivariate system.⁶

An even more important novelty in our analysis is that we identify three features of the co-movement of various forecasted series that are inconsistent with our basic stochastic growth model. The first is that predictable movements in output and predictable movements in the average product of labor bear little relation to each other and, when they are related to each other, they are often negatively correlated. By contrast, the model associates predictable increases in output with capital accumulation which raises labor productivity. Thus predictable movements in output should be strongly positively associated with predictable movements in productivity. The second is that predictable increases in output are correlated with predictable increases in the labor input. By contrast, under standard calibrations, the growth model implies that the labor input should be above its long run value when output is below its long run value, and vice versa, so that the two quantities are always expected to move in opposite directions. Finally, as has been pointed out repeatedly (see e.g., Campbell 1987, Cochrane 1994a), the data suggest that high levels of the ratio of consumption to output are associated with increases in output (and with smaller increases in consumption). We show that this too is inconsistent with the growth model.

We thus believe that we have identified important respects in which a simple stochastic growth model does not predict aggregate fluctuations of the kind observed. It might be argued that this is a fine point, to be addressed by more sophisticated versions of the RBC model, and that it does not detract from the

⁶The problem that this poses for real business cycle theory is also discussed in Cochrane (1994b).

basic model's success in accounting for the overall variability of output, investment, and so on, and for the correlations between overall variations in these variables. However, the significance that one attaches to the successes typically cited in the RBC literature depends upon how one proposes to separate out the "cyclical" movements in aggregate quantities from the trend growth in those quantities (that most economists would explain in terms of technical progress, regardless of their view of the cycle).

One simple approach to characterization of the cyclical movements that a business cycle theory should explain is to subtract a linear trend from (the logarithm of) each of the aggregate time series, as in King, Plosser and Rebelo (1988a). This approach, however, does not make sense if one accepts the evidence referred to above for the existence of a stochastic trend. An alternative approach, used by many authors following Kydland and Prescott (1982), is to extract a complicated nonlinear trend from each aggregate time series using the "Hodrick-Prescott filter".⁷ This allows for a shifting trend, but has the disadvantage that the shifting trend that is removed from the data (actually, the several shifting trends that are extracted from the different aggregate series) is not modeled; the theoretical model that is used by Kydland and Prescott to explain the cyclical variations in the data involves no trend growth at all, and so the effects of shifts in the trend on the variables' deviations from trend are not modeled. This not only casts doubt upon the reliability of the authors' numerical results; it undermines a principal intellectual appeal of the RBC approach, namely, the prospect of an integrated theory of growth and the cycle.

Authors such as K-P-R and Christiano and Eichenbaum (1992) avoid this dilemma by assuming a random walk in productivity. In this case, the model predicts the existence of a stochastic trend in real activity, and indeed the technology shocks that are to explain short-run fluctuations are nothing other than the shifts in this stochastic trend. There remains, however, the question of how to define "cyclical" variations in the presence of a stochastic trend. These authors (and likewise Plosser (1989)) simply discuss the unconditional moments of the *growth rates* of aggregate quantities such as output and investment, on the ground that both according to the theoretical model and in the U.S. data, the aggregate quantities are non-stationary while their growth rates are stationary (and so have well-defined moments). But it is not obvious that the features of the aggregate data that are emphasized in this way should be taken to characterize the business cycle. Probably the most widely accepted proposal for defining a "cyclical" component of time series of this

⁷For discussion of the properties of this filter and comparison with other methods of detrending, see King and Rebelo (1993).

kind is that of Beveridge and Nelson (1981), namely, to define the “trend” component of a random variable $\{X_t\}$ at date t as ⁸

$$X_t^{\text{trend}} \equiv \lim_{T \rightarrow \infty} E_t[X_{t+T} - T \log \gamma_X] \quad (1)$$

where the constant γ_X is the unconditional expected rate of growth of X . The “cyclical” component is then

$$\begin{aligned} X_t^{\text{cyc}} &\equiv X_t - X_t^{\text{trend}} \\ &= \lim_{T \rightarrow \infty} E_t[X_t - X_{t+T} + T \log \gamma_X] \end{aligned}$$

But in this case, the cyclical component is exactly the *forecastable change* in the variable, over an infinite horizon. Identification of the degree of cyclical variation in various aggregate quantities, and of the co-movements in these cyclical variations, then amounts to the analysis of the forecastable changes in those variables. The only difference in our approach is that we also consider forecastable changes over shorter horizons, and in fact we give greatest emphasis to the forecastable changes over a two- to three-year horizon. One reason for this is that we find evidence in the U.S. data for a particularly high degree of forecastability of output growth over this horizon, and it seems natural to us to identify exactly this phenomenon as “the business cycle”. If one does so, however, one must conclude that the stochastic growth model cannot account at all for either the existence or nature of the cycle.

Of course, a more complex RBC model might do better at explaining the cycle in this sense. In particular, the assumption of a significant forecastable component to productivity growth (instead of a random walk) ought to result in a prediction of much more forecastable output growth as well. We discuss some simple alternative models of technical progress in section 6 below, but do not intend to consider this problem in any generality. Instead, our main point is that a model with stochastic growth *need not* possess a business cycle to any significant extent, and insofar as one does, the features of the model that account for the cycle will be largely *independent* of those that account for the stochastic trend. This point remains valid whether the cycle is ultimately due to *transitory* exogenous shocks to productivity, or to shocks of some other kind. Thus the mere fact that stochastic productivity growth can be inferred from the presence of a stochastic trend in output does not, in itself, provide support for the real business cycle hypothesis as to the source of the cycle.

⁸Here we assume that $\{\Delta X_t\}$ is stationary.

In section 1, we document the forecastable changes in output and other aggregate quantities for the postwar U.S., using a simple vector autoregression (VAR) framework. In addition to showing the importance of these forecastable changes, we show that a definition of “the business cycle” in terms of variations in forecasted private output growth coincides empirically with other familiar definitions; for example, we show that our dating of cycles on these grounds would be similar to the NBER’s, or to what one would obtain from a simple linear detrending of (the logarithm of) private output. In section 2, we review the predictions of a simple stochastic growth model (the real business cycle model of K-P-R) regarding forecastable changes in aggregate quantities. In section 3, we present the numerical predictions from a calibrated version of the model (intended to match the U.S. economy) and compare them to our empirical results. Section 4 considers the effect of varying the preference parameters of the model. We show that this can help the model explain certain features of the forecastable components of output, consumption and hours but only at the expense of more counterfactual implications concerning the unforecastable movements in these variables. Section 5 shows that our conclusions are robust both with respect to transitory measurement error in various series and to the presence of certain other types of transitory disturbances. Section 6 briefly considers alternative stochastic processes for technology, and Section 7 concludes.

2 Forecastable Movements in Output, Consumption and Hours

In this section we describe the statistical properties of aggregate U.S. output, consumption, and hours, with particular attention to the existence of a “business cycle” in the sense of forecastable changes in these variables. We use a three-variable VAR to characterize these forecastable movements.

Our results obviously depend upon the VAR specification used, and so we discuss our reasons for particular interest in the system that we estimate. First of all, we wish to compare the properties of the U.S. data to the predictions of a standard stochastic growth model. This requires that we use series that represent empirical correlates of variables that are determined in that model. The RBC literature has stressed its predictions for the movements in aggregate output, consumption, and hours. Moreover, the estimation of a joint stochastic process for these variables also implies processes for labor productivity (output per hour) and investment (output that is not consumed). Thus, we are in fact estimating the joint behavior of all of the variables with which the model is concerned.

We do not include in our forecasting regressions certain variables, such as interest-rate spreads, that have been argued by Friedman and Kuttner (1992) and others to be useful in forecasting output. The reason is that the co-movement of these variables with output is not modeled in the standard growth model. Of course, it might be considered a defect of the model that it does not explain relationships of that kind, if they are in fact found useful in forecasting. But here we wish to consider the extent to which an RBC model correctly explains the co-movements of the variables that it clearly aims to explain. Also, the variables that we study here are ones that must be modeled in any business cycle theory, so that our characterization of the data should prove useful in the evaluation of a wide range of potential theories.

Secondly, we wish to avoid the contrasting pitfalls of underestimating the forecastability of aggregate fluctuations due to omission of useful forecasting variables on the one hand, and of overestimating their forecastability due to insufficient degrees of freedom on the other. To avoid overfitting, we use a small VAR with only a few lags. On the other hand, the few variables used are ones that are expected to be useful in the identification of forecastable output growth. Because we are interested in forecastable output growth, output growth itself must be one of the variables in our system.

It follows from simple “permanent income hypothesis” considerations that the consumption share in output should forecast future output growth, and the usefulness of the consumption share as a forecasting variable has been verified by Campbell (1987), Cochrane and Sbordone (1988), Cochrane (1994a), and King, Plosser, Stock and Watson (1991) [hereafter, K-P-S-W]. Likewise, the idea that variations in the labor input can be used to predict future changes in output has been used to identify temporary output fluctuations in a VAR framework by Blanchard and Quah (1989) and Evans (1989).⁹ Furthermore, as we explain in the next section, the stochastic growth model implies that expected growth is a function of a certain state variable (the aggregate capital stock deflated by the labor force and the state of technology). According to that model, both the consumption-output ratio and hours relative to the labor force should also be functions of that state variable, and hence either variable should supply all of the information that is relevant for forecasting future output growth.

We still face a choice between several possible measures of output, consumption and hours. One issue

⁹These authors use the unemployment rate, rather than hours, as their measure of variations in the labor input. For our purposes, hours are preferable, because of their clearer relation to the labor input with which the RBC model is concerned. Because detrended private hours are stationary, as discussed below, they can serve as a cyclical indicator in a way similar to the unemployment rate.

is how to deal with hours and output purchased by the government, given that versions of the model used in the RBC literature generally ignore government altogether. We have chosen to interpret the standard model as a model of fluctuations in *private* output and hours; we explain in the next section how such an interpretation is possible even when the government does hire some hours and purchase some of private output.¹⁰ As a consequence, our output measure is real private value added,¹¹ and our hours measure is hours worked in the private sector.¹²

Another issue is the choice of a measure of consumption. Because the consumption decision modeled in the standard growth model is demand for a non-durable consumption good, we use consumer expenditure on non-durables and services as our measure of consumption. This is also the consumption measure that one has the most reason to expect to forecast future output on permanent-income grounds, and it is the one used in the studies of output forecastability mentioned above. This choice has the consequence that we identify consumer durables purchases as part of “investment” in the growth model.

A further issue is how to deal with growth of the labor force (also typically ignored in the literature). In the next section, we construct a model with deterministic growth in the available labor force and show that it implies that hours must be trend stationary. This implication is borne out by the data. The last column of Table 1 reports a rejection, using a Dickey-Fuller test, of the hypothesis that private hours have a unit root once one allows for a deterministic trend.¹³

The time series that we use, then, are the logarithms of private output, consumption of nondurables

¹⁰By this we do not mean to assert that we present a model in which government purchases have no effect upon the determination of private aggregates; our model assumes that government purchases, while non-trivial in size, are non-stochastic. While this specification would obviously be contradicted by the data on government purchases, we cannot model them in a more realistic way without considering a more complex model – in particular, a model with multiple stochastic disturbances to the economy. A similar caveat applies to our treatment of population growth.

¹¹We measure real private value added, or “private output”, as the difference between real GNP and government sector value-added, both measured in 1982 dollars.

¹²Our measure of “private hours” is total man-hours reported by non-agricultural establishments, minus man-hours employed by the government. We are thus implicitly assuming that changes in agricultural hours are proportional to changes in private non-agricultural hours. Because our sole concern below is with variations in the logarithm of private hours, the multiplication of our measure by a constant factor greater than one, to reflect the presence of agricultural hours, would not affect any of our results.

¹³Shapiro and Watson (1988) show, by contrast, that one cannot reject the hypothesis that total hours are non-stationary when one does not control for a deterministic trend. Given the growth in both population and in labor force participation, their finding is not surprising. K-P-R show instead that per capita hours are stationary. We do not use per capita hours for two reasons. The first is that per capita hours have a slight deterministic trend of their own which is probably due to varying participation rates; that is, if one allows for a trend in an autoregression of per capita hours, one can reject the hypothesis of a zero coefficient on the trend (even though the series passes some tests of stationarity, and the estimated trend growth rate is small). Second, the use of per capita variables would require, for consistency, that population growth enter as a state variable of our theoretical model. In this case, stochastic variations in population growth would become a second source of disturbances to the model, in addition to stochastic technical progress. This conclusion would be avoided only if we were to assume deterministic population growth, in which case the use of detrended hours and per capita hours would be equally appropriate.

and services, and detrended private hours. Because we use lower case letters to denote the logarithm of the respective upper case letters, these variables are denoted by y_t , c_t and h_t respectively.¹⁴ As K-P-R emphasize, a standard growth model with a random walk in technology implies that y_t and c_t should be difference-stationary, and these two variables should be co-integrated, since $c_t - y_t$ is predicted to be a stationary variable. Table 1 shows that our data are consistent with these predictions as well. One can reject the hypothesis that the consumption share and the rate of growth of private output have unit roots at the 1% significance level. The difference-stationarity of consumption and output, and the stationarity of the consumption-output ratio, are also reported in K-P-S-W, where the issue is discussed in more detail.

Hence our VAR specification is¹⁵

$$z_t = Az_{t-1} + \epsilon_t \quad (2)$$

where

$$z_t = \begin{pmatrix} \Delta y_t \\ (c_t - y_t) \\ h_t \\ \Delta y_{t-1} \\ (c_{t-1} - y_{t-1}) \\ h_{t-1} \end{pmatrix} \quad \text{and} \quad \epsilon_t = \begin{pmatrix} \epsilon_t^1 \\ \epsilon_t^2 \\ \epsilon_t^3 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

and only the first three rows of A need to be estimated. This autoregression includes only two lags. One reason for ignoring further lags is that, when we included them, these were generally not statistically significantly different from zero.¹⁶ A second reason is that we want to avoid overfitting our VAR. Overfitting is a particular concern in that it could lead us to overstate the extent to which aggregate variables are forecastable, and thus, the extent to which they are subject to cyclical movements. Table 1 also presents the estimates from our VAR. As can be seen from the Table, most parameters are statistically different from zero.

We now turn to the characterization of aggregate fluctuations that can be obtained from the joint stochastic process for these three time series. In Table 2, we present the estimated values for several unconditional second moments of the data, of the kind that have been emphasized in the RBC literature. We focus on the behavior of private output y_t , consumption c_t , investment i_t , detrended private hours h_t , and labor

¹⁴We reserve the notation L_t for private hours, prior to detrending.

¹⁵Note that a VAR of this sort is equivalent, except in the way that lags are truncated, to an error-correction model of the kind estimated by K-P-S-W and by Cochrane (1994a, 1994b).

¹⁶Adding either just a third or both a third and a fourth lag of all variables leads to just two coefficients that are statistically significant at the 5% level: the third lag of detrended hours is a significant explainer of both the growth in output and of the consumption share.

productivity p_t . We have discussed above the measures that we use for output, consumption and hours; the other two series are constructed from these, to ensure that our data satisfy the accounting identities linking these quantities in the theoretical model. Thus we construct a series for the growth rate of investment using only our data on the growth in output and consumption, using the relation

$$s_C \Delta c_t + (1 - s_C) \Delta i_t = \Delta y_t$$

In the next section we derive this accounting relation for our theoretical model, from explicit assumptions about government behavior. To compute ΔI using this formula, we need to know s_C . The ratio $s_C/(1-s_C)$ is the average ratio of consumption to investment spending. Using postwar U.S. data, and letting consumption be equal to consumer expenditure of nondurables and services while investment equals the sum of gross fixed investment and consumer spending on durables, we obtain an s_C equal to 0.70. We similarly construct our series for growth in labor productivity from our series for growth of output and hours, using the relation

$$\Delta p_t = \Delta y_t - \Delta h_t$$

(This follows, up to a constant, from the standard definition of average labor productivity.) The first column of Table 2 presents the unconditional standard deviations of the (stationary) growth rate of private output, the ratios of the unconditional standard deviations of the growth rates of each of the other aggregates just mentioned to the standard deviation of output growth, and a number of unconditional correlations among these series. As is emphasized in the RBC literature, investment growth is more volatile than output growth while consumption growth is less so; productivity growth exhibits considerable variability; and the growth rates of consumption, investment, hours and productivity are all strongly positively correlated with growth in output. (The second column presents the theoretical predictions of the calibrated stochastic growth model discussed in the next section.)

We next compute the expected changes in each of these aggregates implied by our VAR model. We denote the difference between the expected value at time t of y_{t+k} and the current value of y_t by $\widehat{\Delta y}_t^k$. This is given by

$$\widehat{\Delta y}_t^k = B_y^k Z_t \quad B_y^k \equiv e_1'(A + A^2 + A^3 + \dots + A^k) \quad (3)$$

where e_1 is a vector that has a one in the first position and zeros in all others. For the case where $k = \infty$,

we have (minus) the Beveridge-Nelson definition of the cyclical component of $\log Y_t$, which is given by

$$\widehat{\Delta y_t}^\infty = e_1'(I - A)^{-1} A Z_t$$

The expected percent change in consumption, $\widehat{\Delta C_t}^k$ is similarly given by

$$\widehat{\Delta c_t}^k = \widehat{\Delta y_t}^k + e_2' A^k Z_t - (c_t - y_t) \equiv B_C^k Z_t \quad (4)$$

where e_2 is a vector whose second element equals one while the others equal zero and the second equality defines B_C^k . The expected percent change in hours is given by

$$\widehat{\Delta h_t}^k = e_3' A^k Z_t - h_t \equiv B_h^k \quad (5)$$

where e_3 is defined analogously to e_1 and e_2 while the second equality defines B_h^k . The expected percentage changes in investment and in productivity are then computed as linear combinations of these.

Letting V denote the variance-covariance matrix of Z , the variance of $\widehat{\Delta y}^k$ is

$$B_y^{k'} V B_y^k \quad (6)$$

and similarly for the variances of the expected changes in each of the other aggregates. The standard deviations for these expected changes are presented in Table 3. This table also presents a measure of uncertainty for these standard errors. This measure of uncertainty is the standard deviation of the estimate based on the uncertainty concerning the elements of A .¹⁷

The table shows that the standard error of the expected changes for output grows as the horizon lengthens from one to twelve quarters. The largest predictable movements occur in the next twelve quarters and the standard deviation of these predictable movements is above 3.2%. This number can be compared to that in the next to last column, which gives the standard deviation of the unpredictable movements in output over the same horizon. This standard deviation equals only 3.0%. Thus the size of the predictable movements over the next three years exceeds the size of the unpredictable movements. This fact is reflected in the last column, which gives the ratio of the variance of the expected changes over the total variance of output changes. This measure of R^2 is even slightly higher (equal to .55) at the 8 quarter horizon, due to the lower variance of the unpredictable movements at the shorter horizon.

¹⁷We computed the vector of derivatives D of each standard deviation with respect to the elements of A . The variance of our estimate is then $D' \Omega D$ where Ω is the variance-covariance matrix of the elements of A .

For horizons larger than 12 quarters, both the R^2 and, more surprisingly, the size of expected output changes falls. However, the declines in the size of this predictable components are very small and, indeed, are not statistically different from zero given the uncertainty about the parameters of A .¹⁸ On the other hand, the difference between the predictability of output growth at short horizons and the predictability of output growth over the next two to three years is both substantial and statistically significant. The high frequency movements in output are largely unpredictable: less than a third of the variance of output is predictable over the next quarter. By contrast, output movements over two to three years are dominated by predictable “cyclical” components which are our main focus of attention in this paper.

Since the size of these predictable movements is largest at the 12 quarter horizon, we focus mostly on this horizon. (We also prefer this to a longer horizon because the predictable movements over a shorter horizon can be estimated with greater precision.) However, it is important to realize that the expected movements in output over the next eight, twelve or infinite quarters are very similar to each other. To see this, Figure 1 graphs the demeaned expected *declines* in output over these three horizons using the same scale. We show expected declines, as opposed to expected increases, because recessions ought to be associated with expected increases in output and we wish to represent these as low values for our cyclical indicator. In this Figure we have also indicated the troughs of recessions as determined by the NBER. We see that output is expected to grow fast at these NBER troughs so that expected output declines have some similarity with this particular business cycle indicator.¹⁹ It is worth noting, however, that our cyclical indicator tends to reach its lowest value one quarter after the NBER troughs. The reason may be that the NBER uses subsequent information to construct its chronology of troughs. Thus the end of the recession is defined to occur just before output grows more than it was expected to. This positive innovation in output cannot be predicted with our method. On the other hand, this positive innovation tends to raise predicted output growth since lagged output growth predicts future output growth to some extent.

In Figure 2, we superimpose our measure of expected declines over the next twelve quarters with linearly detrended values of our output measure itself. The Figure shows that our cyclical indicator is very closely associated with detrended private value added. There are some interesting differences, however. First, as

¹⁸ Thus a possible interpretation of our findings would be that output is expected to have essentially completed the adjustment to its long-run value within a period of two to three years, after which little further change in output can be forecasted.

¹⁹ The one case where the indicators differ is in the case of the last recession. As would be suggested by our series, the recovery from this “trough” was initially weak.

with the NBER troughs, detrended output appears to lead slightly our business cycle indicator. Once again, this may be due to the fact that unpredictable upturns in activity are responsible for the turning points of the detrended output series.²⁰

The two lines also differ in some of their low frequency aspects. First, our predicted output growth series is relatively smaller in the 1960's than the deviation from a linear trend. While the linear detrending method attributes all the unusually large growth in the 1960's to an abnormally large cyclical expansion, our method attributes some of it to variables that affect steady state output. By the same token, linearly detrended output falls more in the 1970's than our series. Finally, in the 1980's linearly detrended output remains low in part because the high growth of the 1960's did not persist. By contrast, our series treats the Reagan expansion as unusually large.

In addition to forecastable output movements, we find some forecastable consumption movements as well.

²¹ As Table 3 shows, the volatility of expected consumption growth is substantially smaller than the volatility of expected output growth. This is not very surprising given that the volatility of overall consumption growth is smaller than the volatility of output growth. What is notable about the behavior of consumption is that the relative variability of consumption is particularly small at short horizons. Expected consumption changes are less than half as large as output changes at horizons of under one year. As the horizon lengthens, the standard deviation of expected consumption changes keeps growing while that of output does not. The result is that the standard deviation of the difference between expected *long run* consumption and current consumption is about 88% as large as the corresponding standard deviation for output. Table 3 also shows that expected investment growth is more volatile than expected growth in output. This is not surprising given that expected growth in consumption is less volatile than expected growth in output. The magnitude of the expected changes in hours worked reported in Table 3 is very similar to that of expected changes in output. This may seem surprising given that K-P-R report that the standard deviation of the growth in hours is only 80% as large as the standard deviation of the growth in output. Finally, Table 3 reports the

²⁰From a purely mechanical point of view the finding that our indicator lags behind output is not surprising since our indicator is heavily influenced by hours worked, which are known to lag behind aggregate activity. This raises the question of whether our cyclical indicator would have been different if we had allowed detrended output to influence it. To check this we ran a regression explaining output growth with our variables and, in addition, one lag of detrended output. The coefficient on detrended output was statistically insignificant and the other coefficient estimates did not change. Thus hours, the ratio of consumption to output, and the rate of growth of output contain substantially more information about future output growth.

²¹The existence of such movements – a violation of a simple version of the rational-expectations permanent income hypothesis – has been demonstrated before (e.g., Campbell and Mankiw, 1989).

volatility of expected changes in labor productivity. At horizons above one year, the standard deviation of this is less than one third of the standard deviation of expected output growth.

In Tables 4 and 5 we give further statistics that describe the behavior of cyclical output, consumption and hours indicated by our VAR model. Table 4 is devoted to studying which of our regressors is particularly responsible for forecasting output growth at various horizons. Table 5 focuses instead on the expected co-movements of the five aggregate variables.

Table 4 gives the correlations of $\widehat{\Delta y}_t^k$ and the value at t of our three regressors, Δy_t , $(c_t - y_t)$ and h_t . The forecast of output growth in the next quarter is most highly correlated with current output growth. By contrast, the other two variables, especially detrended hours, are more useful for predicting output growth over longer horizons. Thus hours prove to be an important indicator of the current state of the business cycle.

Table 5 presents regression coefficients of the expected changes in c , h , p and i on expected changes in y . These coefficients indicate the percentage by which a given variable can be expected to change if one knows that output is expected to increase by one percent. The covariance between expected consumption growth over the next k quarters and expected output growth over this same interval is given by $B_c^{k'} V B_y^k$. Thus, the regression coefficient relating changes in c at horizon k to the changes in y over the same period is given by

$$\frac{B_c^{k'} V B_y^k}{B_y^{k'} V B_y^k} \quad (7)$$

The other coefficients can be computed analogously.

Table 5 indicates that the elasticity of expected consumption growth with respect to expected output growth equals about one-fourth for a one quarter horizon. It grows with the horizon so that it exceeds one-half for the infinite horizon. Given that the standard deviation of consumption changes is smaller than that for output, it is not surprising that consumption does not respond one for one to expected changes in output. While expected consumption growth is not very elastic with respect to expected output growth, investment growth is, and this is consistent with the large volatility of investment changes.

Expected hours growth responds nearly one for one to expected changes in output. As a result, expected productivity is largely unrelated to expected changes in output. This lack of correlation may be surprising given that output growth is generally positively correlated with productivity growth. The table indicates

that this correlation is due mainly to unexpected changes in either output or productivity. The data are thus consistent with the idea that measured productivity growth is strongly associated with current shocks.

The regression coefficients in Table 5 can be used together with the standard deviations reported in Table 3 to compute correlations between expected changes in consumption and hours on the one hand and expected changes in income on the other. Particularly for horizons above 8 quarters the regression coefficients of both consumption and hours are close to the ratio of their respective standard deviations to that of expected output growth. If they were equal, $\widehat{\Delta y}^k$ and $\widehat{\Delta c}^k$ (or $\widehat{\Delta y}^k$ and $\widehat{\Delta h}^k$) would be perfectly correlated. As it is, the correlation between expected changes is high but not equal to one. The correlation between expected consumption growth over the next 8 quarters and expected output growth is .79 while that between expected hours growth and expected output growth is 0.97. The high values of these correlations suggest that there is a single underlying state variable that determines the position of the economy in the business cycle and hence the evolution of expected output, consumption and hours.

In the stochastic growth model that we consider in the next section, there is a state variable of this sort. As we show, this state variable is the ratio of the current capital stock to the capital stock that is expected to exist in the infinite future. The question then becomes whether this state variable can explain the size and nature of the movements and co-movements documented in this section.

3 A Simple Stochastic Growth Model

In this section we describe the structure of a stochastic growth model, the predictions of which we wish to compare with the properties of the aggregate time series just discussed. The model extends the stochastic growth model of Brock and Mirman (1972) to allow for a labor-leisure choice; it is essentially identical to the model analyzed in K-P-R and in Plosser (1989). We consider this variant, rather than familiar alternatives such as the models analyzed in Prescott (1986), because it implies a stochastic trend of the kind assumed in our treatment of the data. The primary innovations in our own presentation of the model are explicit treatment of government purchases and labor force growth, in order to tighten the relation between the theoretical model and our time series.²²

²²The predictions of the model that we present here are in fact almost identical to those of the K-P-R model. We contrive our treatment of the government so that our model's predictions regarding the evolution of (private) capital, hours, consumption, and output are identical to those of a model with no government.

Consider an economy made up of a fixed number of identical infinite-lived households. We will suppose that there is growth over time in an exogenous state variable N_t that we refer to as the size of the labor force,²³ but that in our model simply represents a change in the preferences of the representative household. (Each household may be supposed to be made up of many individuals whose number may grow with time.) The representative household seeks to maximize the expected value of lifetime utility

$$U = E_0 \sum_{t=0}^{\infty} \beta^t u \left(C_t, \frac{L_t^{\text{tot}}}{N_t} \right)$$

where β is a constant discount factor, C_t again denotes consumption by the members of the household in period t , and L_t^{tot} denotes per total hours worked by the members of the household in period t . We let $H_t^{\text{tot}} \equiv L_t^{\text{tot}}/N_t$ be total hours worked as a ratio to the available labor force. The single-period utility function $u(C, H)$ is concave, increasing in C , and decreasing in H .

In order to ensure the existence of a stationary equilibrium (in terms of suitably rescaled state variables) despite the presence of a non-stationary technology process (specified below), we need further homogeneity assumptions on the form of this function. The marginal utility of income λ_t for the representative household each period must be given by

$$\lambda_t = u_C(C_t, H_t^{\text{tot}}) \quad (8)$$

Furthermore, household optimization requires that each period the marginal rate of substitution between C and L^{tot} must equal the real wage W_t , so that

$$\frac{u_h(C_t, H_t^{\text{tot}})}{u_C(C_t, H_t^{\text{tot}})} = -W_t N_t \quad (9)$$

Equations (8) and (9) then implicitly define Frisch labor supply and consumption demand functions $H^{\text{tot}}(W_t N_t, \lambda_t)$ and $C(W_t N_t, \lambda_t)$, which provide a useful characterization of intra-temporal preferences.²⁴ Our additional homogeneity assumption can then be stated as follows. There exists a parameter $\sigma > 0$ such that the Frisch labor supply is homogeneous degree zero in $(WN, \lambda^{-1/\sigma})$, and the Frisch consumption demand is homogeneous of degree one in the same arguments.²⁵ The consequence of this assumption is

²³Our “labor force” variable is a scale factor for aggregate labor supply, with no necessary connection with the labor force measure in the BLS surveys, that measures only those individuals who either have a job or who are actively seeking one, and does not weight them according to the amount of time that they wish to work.

²⁴See, e.g., Rotemberg and Woodford (1992, 1994) for further discussion of Frisch labor supply and consumption demand functions.

²⁵The family of utility functions with this property is discussed further in King, Plosser and Rebelo (1988a) and in Rotemberg and Woodford (1992).

that a permanent increase in the real wage leaves hours worked unchanged while desired consumption rises proportionally. An increase in N has the same effects as an increase in w since it increases the payment that the household receives per unit of H^{tot} . Thus an increase in the labor force also lacks any effect on the value of H^{tot} , while C and L^{tot} grow in proportion to the increase in N . Note that the case analyzed in much of the RBC literature,

$$u(C, H) = c - v(H),$$

for $v(H)$ an increasing convex function, satisfies our homogeneity assumption with $\sigma = 1$.

We assume a constant rate of growth of the labor force N_t so that our model has only a single source of stochastic variation in the endogenous variables. Thus $N_t = N_0 \gamma_N^t$ for some positive constants N_0 and γ_N .

Private output is produced by competitive firms using a technology

$$Y_t = F(K_t, z_t L_t)$$

where Y_t denotes private output as before, K_t the private capital stock, L_t is private hours, and z_t an exogenous technology factor, all in period t . Stochastic variations in the technology factor are the source of aggregate fluctuations. We introduce a stochastic trend in output, consumption, and so on (as was argued to exist in the U.S. data) by assuming that the technology factor is a random walk with drift, i.e., that

$$\log z_t = \log \gamma_z + \log z_{t-1} + \epsilon_t \quad (10)$$

where γ_z is a positive constant and $\{\epsilon_t\}$ is a mean-zero i.i.d. random variable.

Both factors of production are hired in competitive spot markets each period. The evolution of the private capital stock is given by

$$K_{t+1} = I_t + (1 - \delta)K_t$$

where I_t denotes private investment and the depreciation rate δ is a positive constant, less than or equal to one.

The government is assumed to hire a certain constant fraction of the available labor force for its own use. We denote this fraction by H^g , it is recorded in the national income accounts as government value added. As a result, the condition for labor market clearing is

$$H^{\text{tot}}(W_t N_t, \lambda_t) = H_t + H^g \quad (11)$$

where H_t denotes private hours per member of the labor force, or L_t/N_t .

In addition, the government purchases a quantity G_t of private output in period t . We assume that

$$G_t = \tau Y_t$$

where τ is a positive constant, less than one. As a result, the condition for product market clearing is

$$C_t + I_t = (1 - \tau)F(K_t, z_t N_t H_t) \quad (12)$$

Both H^g and τ are assumed to be constants (at some cost in realism) so that the technology shock is the only source of stochastic variation.

These government purchases are financed by lump-sum taxes equal to $W_t N_t H^g$ each period, and a proportional tax rate τ on all factor incomes. This results in equilibrium factor demands satisfying

$$\begin{aligned} F_1(K_t, z_t N_t H_t) &= \frac{\rho_t}{1 - \tau} \\ F_2(K_t, z_t N_t H_t) &= \frac{W_t}{(1 - \tau)z_t} \end{aligned}$$

where ρ_t denotes the after tax rental price of capital goods in period t . These equations hold because $\rho/(1 - \tau)$ and $W_t/(1 - \tau)$ equal the pre-tax wage and rental rate respectively. These equilibrium conditions, along with (12), are identical to those of a model with no government purchases in which the production function $F(K, zL)$ is replaced by $(1 - \tau)F(K, zL)$ and the Frisch labor supply function $H^{\text{tot}}(W_t N_t, \lambda_t)$ is replaced by $H^{\text{tot}}(W_t N_t, \lambda_t) - H^g$.²⁶ Because neither the steady changes in N_t nor the permanent changes in W_t induced by technology have a permanent effect on H , this variable is stationary. Thus, as we suggested in the earlier section, private hours must be trend stationary.

The equilibrium for this economy is given by the solution to a planning problem. The levels of output, consumption and hours that solve this planning problem at any date t depend on K_t , N_t , the current level of technology at t , z_t and on the expected evolution of technology in the future. Because z_t is Markovian, the expected future evolution of technology depends only on z_t itself so the equilibrium at t depends only on K_t , N_t and z_t . Moreover, it is easy to show that our preference specification implies that, output Y_t

²⁶Note that this does not involve any violation of the usual properties of the Frisch demands; the modified functions $H(W_t N_t, \lambda_t)$ and $C(W_t N_t, \lambda_t)$ are simply the Frisch demands corresponding to a modified utility function $u(C, H + H^g)$.

and consumption C_t (and hence investment) are homogeneous of degree one in K_t and $z_t N_t$, while H is homogeneous of degree zero in these same two variables.²⁷ This means that the rescaled levels of private consumption, investment, output and hours at date t , $\frac{C_t}{z_t N_t}$, $\frac{I_t}{z_t N_t}$, $\frac{Y_t}{z_t N_t}$, and $\frac{L_t}{N_t}$ respectively, are each functions of just $\frac{K_t}{z_t N_t}$, and we denote the logarithm of this state variable by κ_t . As a result of this, the rescaled level of labor productivity $P_t/z_t \equiv \frac{Y_t}{z_t L_t}$ is also a function of κ_t . It can be shown furthermore that κ_t is a stationary variable in the equilibrium, and hence that each of the rescaled variables just mentioned is stationary as well.

In a log-linear approximation to the equilibrium laws of motion, we can then write

$$(x_t - x^*) = \pi_{x\kappa}(\kappa_t - \kappa^*) \quad (13)$$

where x_t denotes the logarithm of any one of the five stationary variables just mentioned, x^* is the mean value of that variable, and κ^* is the mean value of κ_t . (We will use subscripts c, i, y, h and p for x in referring subsequently to these elasticities.) The corresponding investment equation, together with the random walk in technology, then implies a law of motion for the state variable $\{\kappa_t\}$, an approximation to which is

$$(\kappa_{t+1} - \kappa^*) = \eta(\kappa_t - \kappa^*) - \epsilon_{t+1} \quad (14)$$

where $\eta \equiv (1 - \theta)\pi_{i\kappa} + \theta$, and $\theta \equiv (1 - \delta)/\gamma_z \gamma_N$ is the average fraction of the capital stock made up of undepreciated capital from the previous period (as opposed to investment purchases during the previous period). For the calibrated parameter values discussed in the next section, $0 < \eta < 1$, so that (14) implies that $\{\kappa_t\}$ is indeed a stationary variable. Given an initial per-capita capital stock k_0 and an initial state of technology z_0 , equations (13) and (14) determine the evolution of the variables $\{z_t, K_t, C_t, I_t, Y_t, h_t, P_t\}$ as a function of the sequence of technology innovations $\{\epsilon_t\}$.

These equations can thus be used to compute impulse response functions to a disturbance ϵ_t . As is apparent from (14), this response simply describes the expected evolution of our variables starting from a situation where κ_t is away from its steady state value κ^* . This evolution is thus identical to the deterministic dynamics of a model with constant technology whose initial capital stock is different from its steady-state level.²⁸ Equation (13) then determines the extent to which each particular variable departs from the steady

²⁷ See K-P-R for discussion of this, and also the log-linear approximation to the equilibrium dynamics used below.

²⁸ This problem is analyzed in section 3 of King, Plosser and Rebelo (1988a); their notation for the elasticity η is μ_1 .

state while (14) implies that they all converge exponentially to the steady state at the rate $(1 - \eta)$. The fact that $0 < \eta < 1$ indicates that the system does indeed converge. For plausible parameter values (in particular, a plausibly low depreciation rate δ), the rate of convergence is relatively slow, so that η is near 1.

Using the baseline parameters set out in Table 6 and discussed further below, Figure 3 displays the response of consumption, hours and output to a disturbance that raises ϵ by one. From an initial value of zero, the increase in ϵ eventually raises output and consumption by one unit while hours return to their original value. Since the increase in ϵ lowers the capital stock relative to its steady state value, output is below the new steady state as well. These parameters imply that when output is below the steady state, consumption is even further below the steady state. This ensures that the ratio of investment to output is above the steady state and helps raise the capital stock to its steady state value. The figure also shows that hours are above the steady state. This occurs because the low value of the capital stock implies that wealth is relatively low so that people reduce their consumption of leisure.

The joint stochastic process for these variables is predicted to be such that Δy_t , $(c_t - y_t)$, and h_t are stationary variables, though $\{c_t, y_t\}$ each possess a unit root, as is reported in section 2 for the U.S. data. Specifically, the model predicts that

$$\Delta y_t = \pi_{y\kappa} \Delta \kappa_t + \epsilon_t \quad (15)$$

$$(c_t - y_t) = [\pi_{c\kappa} - \pi_{y\kappa}] \kappa_t \quad (16)$$

$$h_t^p = \pi_{h\kappa} \kappa_t \quad (17)$$

omitting the constants in each equation. On the other hand, both c_t and y_t are non-stationary, as each can be expressed as the sum of $\log z_t$ (a random walk) and a stationary variable. Hence the general form of our econometric specification in section 2 is consistent with this model.²⁹

K-P-R describe the numerical predictions of the model regarding the variability of the *growth rates* of aggregates such as per capita output, consumption, investment and hours. We wish to emphasize instead the character of predictable changes in such aggregates. In the case of the aggregates $X_t = C_t, I_t, Y_t$, or P_t ,

²⁹The model does imply that the joint stochastic process for the three stationary variables should be singular, as there is only a single shock each period to which all three innovations are proportional, and this is not true of the VAR that we estimate in section 2. But this is not a prediction of the model that we propose to test here; in this particular respect it is obvious that a one-shock model is inadequate. It is still of interest to ask to what extent a particular one-shock model predicts co-movements of aggregate variables that are at all similar to those observed. If it does, there might be some hope that one shock could be responsible for the greater part of what are thought of as typical “business cycles”.

the laws of motion (13) and (14) imply that the expected growth rates are given by

$$\begin{aligned}\widehat{\Delta x}_t^k &= \pi_{x\kappa} E_t[\kappa_{t+k} - \kappa_t] + E_t[\log z_{t+k} - \log z_t] \\ &= -\pi_{x\kappa}(1 - \eta^k)(\kappa_t - \kappa^*)\end{aligned}$$

In writing this, it is assumed that the date t information set used in forecasting includes the state variable κ_t . However, since both $(c_t - y_t)$ and h_t^p are log-linear functions of this variable, it suffices that either of the latter variables be in the information set. Since both of these variables are among our regressors in section 2, the model implies that the forecastable changes identified by our VAR specification should coincide with the variables described above. Similarly, the laws of motion imply that

$$\begin{aligned}\widehat{\Delta h}_t^k &= \pi_{h\kappa} E_t[\kappa_{t+k} - \kappa_t] \\ &= -\pi_{h\kappa}(1 - \eta^k)(\kappa_t - \kappa^*)\end{aligned}$$

Thus the forecastable changes in all of our five variables are predicted to be perfectly collinear. Furthermore, the forecastable changes in any one of the variables at different horizons are predicted to be perfectly collinear: for regardless of the horizon k , the forecastable change should be proportional to $(\kappa_t - \kappa^*)$.

4 Numerical Results for the Baseline Model

We now present the numerical predictions of a calibrated version of the stochastic growth model described in the previous section, and compare them to our estimates in section 2. The calibrated parameters presented in Table 6 are identical to those used by K-P-R, except that we allow for growth in the labor force. A regression of the logarithm of private hours on a deterministic trend gives γ_N , the rate of growth of the labor force. Our regression implies that this equals 1.004.

Note that preferences are specified in terms of the coefficient σ referred to in the previous section – that can be interpreted as the reciprocal of the intertemporal elasticity of substitution of consumption holding hours worked constant – and ϵ_{HW} , the elasticity of the Frisch labor supply function $h^{tot}(wN, \lambda)$ with respect to the real wage. As is explained in Rotemberg and Woodford (1992, 1994), the other elasticities of the Frisch demands can all be computed given numerical values for these two parameters, and the elasticities of the Frisch demands are the only aspect of preferences that enters the log-linearized equilibrium conditions.

The values used in our baseline calibration – $\sigma = 1, \epsilon_{HW} = 4$ – are those that would result from a utility function

$$\log(C_t) + \log(\bar{H} - H_t^{\text{tot}})$$

if on average H^{tot} is .2 of $\bar{H} - H^g$.

The other parameter not taken directly from K-P-R is our assumed standard deviation for the technology shocks, $\sigma_\epsilon = .00732$. This value is equal to the estimated standard deviation of innovations in the permanent component of private output, from the VAR described in section 2.³⁰ According to the theoretical model of the previous section, the trend component of log private output in the sense of (1) should exactly equal $\log z_t$ (plus a constant), so that the variance of innovations in this variable should equal the variance of $\{\epsilon_t\}$.³¹ Thus we calibrate the variability of the innovations in technology so that the model's prediction regarding the variability of the permanent component of private output agrees exactly with what we measure. Of course, the fact that the model predicts variation in the permanent component does not imply anything about variation in forecastable changes in output; for example, if output were predicted to be a random walk, there would be none. We turn next to the model's predictions regarding the size and character of the fluctuations in the aggregate variables discussed in section 2.

First of all, the second column in Table 2 presents the predictions of the calibrated model for each of the unconditional moments reported in that table. This is the type of test of the model emphasized by K-P-R and by Plosser (1989). Using this test, the model meets with a fair degree of success. This picture of relative success changes considerably, however, if one considers the variability of the forecastable changes in the various aggregate variables, rather than the unconditional variability of their growth rates. Table 7 reports the predicted standard deviations of the forecastable changes Δx_t^k , for each of the five variables x , and for several different horizons k . The first thing to notice about these results is that the stochastic growth model does not predict that there should be a great deal of variation in the forecasted change in private

³⁰Our results are in essential agreement with those of K-P-S-W, who report a standard deviation of .007 for the “balanced-growth shock” to their three-variable VAR, which differs from ours mainly in using the share of fixed investment in private output, rather than private hours relative to the labor force, as the third variable.

³¹It has been observed by Lippi and Reichlin (1993) that identification of shifts in the permanent component of output using a VAR in this way depends upon an assumption of “fundamentality” of the moving-average representation implied by the estimated VAR, an assumption that need not be valid in general. That is, it need not be possible to recover the true permanent shock as any linear combination of the VAR innovations. However, in the present case, our theoretical model implies that the MA representation derived from our VAR should indeed be “fundamental”. The true permanent shock can indeed be recovered from the VAR residuals; for example, equations (14) and (17) imply that ϵ_t should be exactly proportional to the residual from a regression of h_t on h_{t-1} .

output. At the 12-quarter horizon, the standard deviation of the forecastable change in output is predicted to be .0029, whereas we estimate it to be .0326 – the model accounts for variations in forecastable output growth of only 9% of the amplitude of the observed variations! At the infinite-horizon (the Beveridge-Nelson cyclical component of output), the model accounts for variations of only 21% of the amplitude of the observed variations.³²

These results depend on assuming that the standard deviation of the technology shocks σ_ϵ is equal to 0.00732. As we explained earlier this is the standard deviation of the shock to the permanent level of output. This standard deviation also implies that the model’s overall standard deviation of output changes is below the actual standard deviation. This can be seen by comparing the 6th column of Table 7 with the corresponding column of Table 3. These columns present the standard deviations of the unexpected changes in output from one period to the next. The standard deviation of unexpected changes over one quarter predicted by the model equals about 60% of the actual one. For the 24 quarter horizon, the model predicts a standard deviation of unexpected changes that is much closer to the actual one.

Obviously, one could raise both the predictable and the unpredictable variability of output changes generated by the model by raising one’s estimate of σ_ϵ . But our value of σ_ϵ is not solely responsible for the results concerning the lack of predictable output changes in the model. To see this, suppose that we set σ_ϵ so that the standard deviation of overall quarterly changes in output predicted by the model equals 0.012, the actual standard deviation of the one quarter changes in private value added. This requires that σ_ϵ be equal 0.0157, which is more than twice as large as our estimate. Even then, the model’s predicted standard deviation of expected output changes over eight quarters equals only 18% of the standard deviation implied by our VAR.

Another way to see that the lack of predictable movements is not solely due to our choice of σ_ϵ is to compare the R^2 ’s in the last columns of the two tables. These R^2 ’s give the ratio of the variance of expected changes to the total variance of changes in output and are thus independent of the level of σ_ϵ . The R^2 ’s predicted by the model are much lower than the actual ones. At the 12 quarter horizon, the estimates imply

³²If one considers instead the fraction of the variance of the forecastable change that is predicted by the model, one finds that the model accounts for only 1% in the former case and only 4% in the latter. This is a preferable metric in certain respects because the total variance equals the sum of the variances induced by independent shocks. Thus, if a successful model could be found that added other independent disturbances to this model, the other shocks would have to account for 99% of the variance of the forecastable changes in the former case, and 96% in the latter.

that over 50% of the variance of output is predictable. By contrast, if the model were correct, only about 2% of the variance of output over this horizon would be predictable.

Another difference between the model and the data is that the model predicts that the standard deviation of forecasted changes in output rises whenever the horizon is lengthened. By contrast, the data suggest that this standard deviation peaks at the 12 quarter horizon, or at any rate increases little after that horizon. Similarly, the model predicts that the R^2 should rise from the 12 to the 24 quarter horizon and this is not true of our empirical results. Thus the forecasted fluctuations predicted by the model have a somewhat different character than those we find in the data. The model's forecasted changes involve adjustments to the steady state that occur over very long spans of time. Instead, the data suggest that the large forecastable changes occur over shorter "business cycle" frequencies. This finding is related to the demonstration by Watson (1993) that the model is unable to replicate the fact that spectra of output growth have a great deal of power at business cycle frequencies.

Perhaps the most counterintuitive contrast between the model and the data concerns the behavior of the variability of consumption. As we saw, the estimated standard deviation of expected consumption changes equals between one third and one half the corresponding standard deviation for output. By contrast, the model predicts that the standard deviation of expected consumption changes should equal over twice the standard deviation of output changes. This may be surprising since the RBC literature often counts the prediction of relatively smooth consumption as one of the model's important successes. But because technology shocks raise the marginal product of capital they raise interest rates and this promotes a reduction in consumption relative to its steady state level. This reduction is so large in the case of our preference parameters that the ratio of consumption to income actually falls. This means that consumption is expected to grow more than income and thus the size of expected changes in consumption exceeds the size of expected output changes. Another way of seeing this is to note that, in Figure 3, departures of output from the steady state are associated with even bigger departures of consumption from its steady state. Thus the predictable movements of consumption (towards its steady state) are larger.

In the case of investment, by contrast, the model is more accurate. While its underprediction of the total variability of forecastable output movements leads it to underpredict the standard deviation of investment movements, it correctly predicts that this standard deviation should be larger than that for output move-

ments. Investment is very large in the immediate aftermath of a positive technology shock because capital is below the steady state. Later, investment is much smaller and, for this reason, the predictable movements are large. In the data the ratio of the standard deviation of investment movements is to that of output movements is actually slightly larger than the ratio predicted by the model. This is just the flip side of the model's relative overprediction of consumption movements.

The model generates predictable movements in hours that are of roughly the same magnitude as the predictable changes in output. This prediction is validated in the data. This is interesting because, as far as the total variability is concerned, Table 2 shows (as do K-P-R) that the model underpredicts the ratio of the standard deviation of hours growth to that of output growth.

Unlike in the case of output movements, the model predicts labor productivity movements that are too large, particularly for horizons longer than 24 months. This means that the ratio of the standard deviation of productivity changes to that of output changes is much larger in the model than in the data, particularly at long horizons. As we will emphasize below, these counterfactual predictions concerning productivity movements are particularly bothersome because it seems unlikely that simple variants of the model can account for it.

Because the model has just one state variable, κ , the expected changes in all the variables are perfectly correlated. Moreover, since κ is deterministically related to both current hours and the consumption share, these variables are also perfectly correlated with all expected changes. Such perfect correlations are obviously absent from our data. Nonetheless, all our estimates of expected changes are highly correlated with each other and, at least for long horizons, they are also highly correlated with initial hours and the initial consumption share. This suggests that a model with a single state variable can in principle explain a large fraction of the cyclical movements in our variables.

While these correlations are high in our data, their sign is often not that predicted by our model. It is apparent from Figure 3 that our parameters imply that when output is below the steady state (and rising), hours are above the steady state while the consumption share is below the steady state. This means that expected future output growth should be positively associated with the current level of hours and negatively associated with the consumption share.³³ Empirically, both these correlations have the opposite sign from

³³Note that this latter implication is the opposite of what is implied by the simple permanent income hypothesis.

that predicted by our model. In the data, a high consumption share and a low level of hours both predict high future growth in output.

Similar difficulties arise when we analyze the sign of the correlations between the expected changes in our variables. These can be seen in Table 8 which reports the regression coefficients of expected changes in various variables on expected output growth. We do not report different coefficients for different horizons because the model predicts that these coefficients are independent of the horizon in question. If the model were literally correct these regressions would have no error but this is not our main concern here. Rather we are most interested in the form of the predicted co-movement.

Our model predicts that expected consumption growth is positive when expected output growth is positive. Therefore, the regression coefficient of expected consumption growth on expected output growth is positive. Our earlier discussion suggests the model predicts that this coefficient is well above 2. By contrast, it is less than .5 in the data. The problem is once again that the model predicts that consumption will rise faster than income after a positive technology shock.

By the same token, the model implies that expected changes in investment are very negatively related to expected increases in output. The corresponding regression coefficient is about -1.7. Investment is highest immediately after a positive technology shock. After that, investment is expected to fall while output is expected to rise. By contrast, our estimated regression coefficient is above 2, and implies that expected investment growth is positively related to expected output growth.

With our parameters, our model predicts that the regression coefficient of expected hours growth on expected output growth is negative. The reason is that a positive technology shock leads to an immediate increase in hours. As is clear in Figure 3, hours are then expected to fall even though output rises as a result of capital accumulation. In the data, expected output growth is positively associated with expected hours growth. This is simply the flip side of the observation that a low level of hours is associated with an expected increase in output in the data while it is associated with a decline in the model.

A somewhat different contrast is provided by the regression coefficient of expected labor productivity growth on expected output growth. The model predicts this to be positive and larger than one. This is not surprising since output is expected to rise when hours are expected to fall. The extra output is expected to be produced by increased capital. By contrast our estimates indicate that expected productivity growth

is nearly unrelated to expected output growth. While the standard errors are large relative to the point estimates, many of the coefficients in the last column of Table 5 are negative suggesting that productivity should fall when output rises.

5 Alternative Preference Specifications

One obvious question that arises at this point is whether these discrepancies can be resolved by changing the preference parameters in plausible ways. To shed some light on this question, we have investigated whether changes in σ and ϵ_{HW} could reverse the sign of some of the predicted correlations in ways that would make them more consistent with the data.³⁴ The results are presented in Table 9. This table presents the standard deviation of output changes forecasted to occur in the next 12 quarters in the third column. The fourth and fifth columns present the correlation of output changes with the initial level of (C/Y) and h respectively. Finally, the last three columns report the regression coefficients of expected consumption growth, expected hours growth and expected productivity growth on expected output growth.

It is apparent from this table that, keeping σ equal to 1, varying ϵ_{HW} does not make the model more successful. As is well known, raising the elasticity of labor supply, raises the immediate increase in hours in response to a positive technology shock. This means that hours are expected to decline further after such a shock. The result is that the regression coefficient of expected hours changes on expected output changes becomes even more negative. Moreover, the volatility of expected output falls because the increase in output due to capital accumulation is now offset by a bigger decline in hours worked. Thus raising the elasticity of labor supply, which has been demonstrated to help the model explain the unconditional volatility of hours makes the predictions concerning the expected changes in hours more counterfactual.

It is possible to make some of the model's predictions more consistent with our facts by changing σ . Consider first the effect of lowering σ so that utility is more nearly linear in consumption. This has the effect of making consumption rise less in response to a technology shock so that C/Y falls by more. The fall in consumption also tends to raise labor supply so that hours rise by more. If, one assumes also that ϵ_{HW} is large, then hours can rise so much that output initially overshoots its long run level. This possibility is

³⁴These parameters of the model differ from the others in that they cannot be identified from observation of the long run averages of stationary variables alone, as opposed to using evidence on the character of aggregate fluctuations to pin them down. The parameter ϵ_{HW} in particular is quite controversial even within the RBC literature, and the consequences of assuming alternative values for it are often considered. See, e.g., Prescott (1986) and King, Plosser and Rebelo (1988a).

illustrated in Figure 4 which shows impulse responses for σ equal to 0.6 and ϵ_{HW} equal to 10.

Because output overshoots its long run level, output and hours are expected to decline together. This expected decline in output is associated with a low initial level of (C/Y) and a high initial level of hours. Thus the model now fits the sign of these correlations as well as the positive co-movement of expected hours and expected output. Another possible advantage of these parameters is that they imply that labor productivity moves in the opposite direction as output. The reason is that the output declines that follow a positive technology shock are accompanied by capital accumulation and labor decumulation, both of which raise labor productivity. The problem, however, is that, because both factor movements are causing labor productivity to rise, the predicted relationship is too strong; productivity is expected to fall by more than 4% for each 1% increase in output.

There is another problem with this specification of preferences. Because consumption is expected to grow after a technology shock, its growth occurs when output is expected to decline (the regression coefficient of expected consumption growth on expected output growth equals -5.3). While the precise magnitude of this coefficient depends on the parameters employed, it should be clear that the fact that consumption and output move in opposite directions is an immediate consequence of making σ so low that output overshoots its long run level. Thus, such low σ 's do not seem appealing.

The alternative is to consider high levels of σ . Raising σ lowers the elasticity of substitution of consumption so that consumption rises more in the immediate aftermath of a technology shock. This has two effects, both of which make the model more consistent with the data. The first is that, if the increase in consumption is large enough, the consumption share actually rises. Since output is still predicted to grow after the shock, the current consumption share becomes positively correlated with expected output growth. The second effect of raising σ is that, because consumption rises more, hours rise less. If one also lowers the elasticity of labor supply, then hours actually fall after a positive shock to technology. This means that hours are expected to rise together with output, as our estimates suggest.

We document these effects in Figure 5 where we plot the impulse response functions when σ is equal to 4 and ϵ_{HW} is equal to 0.2. As Table 9 indicates, these parameters again imply that high values of h and low values of (C/Y) are associated with output declines. They also imply that the predicted elasticity of consumption growth with respect to expected output growth is only .90 and this lower value is more

consistent with our estimates. Indeed, the fact that this elasticity is below $1/s_C$ implies that the elasticity of expected investment growth with respect to expected output growth is positive, as in the data. Finally, the elasticity of expected hours growth with respect to expected output growth is positive although it equals only 0.06 which is much lower than in the data. Nonetheless, these preference parameters capture the main qualitative features of the forecasted movements in consumption, output and hours. The exception is that they still imply that productivity is expected to grow substantially with output, and this is not the case in our data.

While the use of these preference parameters improves the ability of the model to explain the correlations reported in Tables 4 and 5, it worsens its ability to explain several of the moments reported in Table 2. In particular, the predicted standard deviation of the overall one quarter change in hours falls significantly. It now equals only about 3% of the standard deviation of changes in output. On the other hand, the model now predicts an excessive volatility of consumption. The predicted standard deviation of consumption changes from one quarter to the next now exceeds the corresponding standard deviation for output. But perhaps the biggest problem with assuming such a low elasticity of substitution of consumption is that it results in a strong negative correlation between the overall change in hours and the overall change in output. This may be surprising because the correlation between the predicted changes in the two variables is now positive, as in the data. The problem is that the predictable movements remain small relative to the unpredictable movements. And, positive shocks to productivity now lower hours while raising output, contributing to an overall negative correlation between these variables.

Thus, using the parameters that fit better the expected changes implies that the unexpected changes must largely be due to shocks other than technology shocks. Hence simple variation of the preference parameters does not solve the problem posed in section 4.

6 Consequences of Measurement Error and Transitory Disturbances

We now consider whether the difficulties of the standard stochastic growth model can be explained by the simple hypothesis that our time series on consumption, output and hours are subject to classical measurement

error.³⁵ That is, we wish to consider the hypothesis that the equations of section 3 correctly describe the evolution of a vector of true state variables $\tilde{x}_t = (\tilde{y}_t, \tilde{c}_t, \tilde{h}_t)$, but that our data are for a vector

$$x_t = \tilde{x}_t + \nu_t \quad (18)$$

where $\{\nu_t\}$ is vector white noise process, independent of the technology shock process and hence of the variables $\{\tilde{x}_t\}$.

We can also consider simultaneously the consequences of adding various types of transitory shocks to the model set out in section 3. Suppose that the equilibrium conditions of section 3 are correct, except for the presence of a vector of white noise disturbance terms ν_t , that may enter any of the equations involved in the determination of variables at date t . For example, we may allow for preference shocks, so that the Frisch demand functions become $h^{tot}(w_t N_t, \lambda_t, \nu_t)$ and $C(W_t N_t, \lambda_t, \nu_t)$. (We continue to assume the same homogeneity properties as before, for each value of ν_t .) We may allow for stochastic government purchases and fiscal policy, so that

$$\begin{aligned} L_t^g &= H^g(\nu_t) N_t \\ G_t &= \tau(\nu_t) Y_t \end{aligned}$$

Or we may allow for stochastic variation in the rate of depreciation, so that

$$K_{t+1} = I_t + (1 - \delta(\nu_t)) K_t$$

We can also allow ν_t to be an argument of the date t production function, as long as the homogeneity and concavity properties of the function continue to hold for all values of ν_t ; in this case there would be, in effect, both permanent and transitory technology shocks. A more complex possibility (that we cannot develop here in detail) would be to allow for a wedge between the marginal products in date t production and factor prices, or for a wedge between the representative household's marginal rate of substitution between consumption and leisure and the real wage, that depends upon the difference between the logarithm of the nominal price level at date t and what this price level was expected to be at date $t - 1$.³⁶ Here we do not wish to discuss

³⁵It has often been suggested in the RBC literature that important discrepancies between the predictions of the stochastic growth model and the statistical properties of aggregate time series are due to measurement error, especially in the hours series. See, e.g., Prescott (1986) or Christiano and Eichenbaum (1992).

³⁶This could result from nominal wages or prices being fixed a period in advance, or from asymmetric information regarding the current price level as in the model of Lucas (1972). See Cooley and Hansen (1993) for discussion of complete stochastic growth models into which such sources of monetary non-neutrality are introduced.

extension of the model to include nominal price level determination, but it is clear that such a “price level surprise” variable must be white noise, regardless of the nature of the underlying shock that determines it.

In any of these cases, the perturbed equilibrium conditions have a solution in which \mathbf{x}_t can be expressed as a time-invariant function of κ_t and ν_t . For the variables \mathbf{x}_t are determined as before by equilibrium conditions that are the same at all times, except for their dependence upon the current value of κ_t , and now also their dependence upon the current value of ν_t and the distribution of possible future histories of the shocks $\{\nu_{t+j}\}$. But the distribution of possible future histories of the shocks $\{\nu_{t+j}\}$ is the same at all times, because of the assumption that the shocks are white noise. (Actually, it suffices that the vector process $\{\nu_t\}$ be Markovian.) Thus \mathbf{x}_t is determined solely by κ_t and ν_t . Log-linearization around the mean values of the latter two variables then yields

$$(\mathbf{x}_t - \mathbf{x}^*) = \pi_{\mathbf{x}\kappa}(\kappa_t - \kappa^*) + \pi_{\mathbf{x}\nu}(\nu_t - \nu^*) \quad (19)$$

as a generalization of (13), and similarly

$$(\kappa_{t+1} - \kappa^*) = \eta(\kappa_t - \kappa^*) + \eta_\nu \nu_t - \epsilon_{t+1} \quad (20)$$

as a generalization of (14). Furthermore, due the usual certainty-equivalence property of the log-linear equilibrium conditions, the coefficients $\pi_{\mathbf{x}\kappa}$ and η are the same as in the previous model (corresponding to the case of ν_t a constant vector).

Both the model of measurement error (18) and the model with white noise disturbances (19)–(20) have a common set of implications. These are that

$$E_t[h_{t+j}] = \pi_{h\kappa} \eta^{j-1} E_t[\kappa_{t+1} - \kappa^*] \quad (21)$$

$$E_t[c_{t+j} - y_{t+j}] = [\pi_{c\kappa} - \pi_{y\kappa}] \eta^{j-1} E_t[\kappa_{t+1} - \kappa^*] \quad (22)$$

$$E_t[\Delta y_{t+j+1}] = -\pi_{y\kappa} (1 - \eta) \eta^{j-1} E_t[\kappa_{t+1} - \kappa^*] \quad (23)$$

for any $j \geq 1$, where the coefficients $\pi_{\mathbf{x}\kappa}$ and η are the same as in the previous model. Note that we have used the assumption that the disturbances ν_t in (18)–(20) are white noise to eliminate terms of the form $E_t[\nu_{t+j}]$ for $j \geq 1$. Both models similarly imply that

$$\widehat{\Delta \mathbf{x}_{j,t}}^k \equiv E_t[\mathbf{x}_{t+j+k} - \mathbf{x}_{t+j}] = -\pi_{\mathbf{x}\kappa} (1 - \eta^k) \eta^{j-1} E_t[\kappa_{t+1} - \kappa^*] \quad (24)$$

for any $j \geq 1$, and for each of the variables $x = y, c, i, p$, or h . All of equations (21)–(24) are also implications of the model in section 3; the only difference is that corresponding equations with $t + 1$ replaced by t need no longer hold.

Because each of these conditional expectations is a multiple of the same variable $E_t[\kappa_{t+1} - \kappa^*]$, it follows that their relative variances and their correlations are all identical to those predicted by the model of section 3. Hence these implications of the model can be tested, independently of which variables one believes to be most contaminated by measurement error or of which of the types of transitory disturbances one believes are most important.

Up to this point we have considered the effects of measurement error and transitory disturbances that are independently distributed through time so that forecasts of future variables depend only on $E_t[\kappa_{t+1} - \kappa^*]$. A similar analysis applies to the case where the transitory disturbances or the measurement error follow a moving average process of order m . Then, the expectations at t of variables at $t + j$, where j is no smaller than $m + 1$, depend only on the $E_t[\kappa_{t+m+1} - \kappa^*]$. The reason is that any effect beyond $t + m + 1$ of shocks that have impinged on the system up to t must be due only to the slow adjustment of capital from $t + m + 1$ to the steady state.

Table 10 presents the implications of the estimated VAR for several statistics involving expectations at t of k period changes starting at $t + j$. Specifically, for each value of j and k , the first column reports the standard deviation of $\widehat{\Delta y}_{j,t}^k$. The next two columns report the correlations between $\widehat{\Delta y}_{j,t}^k$ on the one hand, and $E_t[c_{t+j} - y_{t+j}]$ and $E_t[h_{t+j}]$ on the other. The remaining three columns report the respective coefficients of the regressions of $\widehat{\Delta c}_{j,t}^k$, $\widehat{\Delta h}_{j,t}^k$, and $\widehat{\Delta p}_{j,t}^k$ on $\widehat{\Delta y}_{j,t}^k$. The theoretical predictions for these last six quantities are independent of both j and k ; the correlations in columns two and three ought to be -1 and +1 respectively (in the case of the calibration described in Table 6), while the regression coefficients in the remaining columns ought to take the values given in Table 8.

Table 10 shows that the co-movements between variables are not very sensitive to the choice of j . The correlations of expected output growth from quarter $t + j$ with the expectation of h at $t + j$ always has the wrong sign as does the regression coefficient of expected hours growth on expected output growth. Because this latter coefficient is always estimated to be near one, expected productivity is either unrelated or negatively related to expected output growth. The sign of the correlation between the expected value

of $(c - y)$ and expected output growth as well as the coefficient in the regression of expected consumption growth on output growth remain inconsistent with the model until j reaches 8 quarters. For longer horizons, these correlations become unstable, though the standard errors become very large as well. The overall stability of the results as one varies the point from which expected changes in output are computed suggests that measurement error and transitory disturbances are not responsible for our results concerning the co-movement of different series.

There remains the question of whether the estimates of the standard deviation of $\widehat{\Delta y_{j,t}}^k$ can be compared to the predictions of the model. In the case of transitory disturbances, such comparison is impossible because, depending on the disturbance, a transitory disturbance at t can have quantitatively important effects on κ_t/κ^* . On the other hand, measurement error at t should not have a large effect on $E_t[\kappa_{t+j} - \kappa^*]$. It is thus of interest to compare the actual variability of $\widehat{\Delta y_{j,t}}^k$ with the variability induced by random walk disturbances with σ_ϵ equal to 0.00732. The corresponding theoretical predictions are displayed in Table 11. We see that, for low values of j , the model still generates predictable movements whose variability is too small. For higher values of j , the correspondence between the two is closer because the data suggest that, beyond a 12 quarter horizon, predictable movements of output tend to be quite small.

7 Slow Diffusion of Technical Progress

One possible answer to the difficulties encountered by the stochastic growth model in the previous sections is to consider a more complex stochastic process for technology. It is rather obvious that one way to obtain larger forecastable movements in output and other variables is to assume forecastable movements in the production technology itself. It is also obvious that generalizing the specification of the technology process can in principle introduce a large number of additional degrees of freedom, perhaps enough to allow one to fit the finite set of statistics discussed above.

But such a resolution, even if possible, does not detract from our main point here. The demonstration above that a particular form of technical progress that yields stochastic growth does not generate business cycles suffices to establish that the business cycle is not a necessary concomitant of stochastic growth. In addition, allowing oneself a large number of free parameters in the assumed technology process – that are not to be pinned down by reference either to microeconomic evidence or to growth facts – is hardly in the spirit of

the RBC literature. This literature has emphasized the benefits of eliminating free parameters whose values are deduced from the cyclical fluctuations that one seeks to explain. It is precisely this desire to conserve on parameters that has led the RBC literature to an almost exclusive focus on the case where z_t follows a random walk. This is the simplest possible specification that allows for stochastic growth; for it is the unique type of process with the property that all conditional expectations of the form $E_t[\log z_{t+j}]$ for $j \geq 0$ – the only aspects of expectations of future technology that matter for equilibrium determination in the log-linear approximation – can be summarized by a single state variable. Furthermore, it has often been argued to be reasonable *a priori* because technical inventions, once discovered, should be permanent additions to knowledge. Thus while one might assume dynamics of some complex sort for expected productivity changes, and thus obtain a model that can explain both a cycle of the kind observed and stochastic growth as consequences of technology shocks, there is no reason to regard such a modification of the model as any less *ad hoc* than would be the introduction of any other additional source of transitory dynamics, such as nominal contracts and monetary policy shocks.

For this reason, we do not here attempt to consider the implications of a general class of stochastic processes for technology; this sort of extension, like the investigation of other sources of transitory dynamics, is left for further work. We do, however, wish to briefly consider an alternative specification for the technology shock process, that is both relatively simple, and that generalizes the random walk specification in a way that is suggested by studies of the actual nature of technical progress, rather than being chosen simply for its usefulness in producing dynamics of the desired sort. Specifically, we wish to allow for the possibility that technical innovations, once discovered, diffuse slowly through the economy rather than being immediately adopted to the fullest possible extent. The evidence that actual innovations are adopted slowly is ubiquitous,³⁷ though the reasons why the diffusion is so slow are unclear,³⁸ and it is not obvious how this specification is to be reconciled with the low serial correlation of measured Solow productivity residuals.³⁹

³⁷See Mansfield (1968) for a comprehensive discussion and for references. Jovanovic and Lach (1993) also discuss the consequences of the slow diffusion of innovations for the ability of a stochastic growth model to account for aggregate fluctuations, although in a growth framework rather different from our own.

³⁸Ellison and Fudenberg (1993) stress the presence of slow learning about the quality of an innovation.

³⁹Under the assumptions of the model analyzed here, the Solow residual should measure the growth rate of the technology factor z_t , so that the observation of serial correlation near zero provides support for the assumption that z_t follows a random walk. However, it is often argued that much of the high-frequency variation in measured Solow residuals results from mis-measurement of inputs (e.g., due to “labor hoarding”, variations in capital utilization, or improper aggregation) rather than true technical progress. Thus the technology specification considered here should probably not be dismissed on this basis alone, though we do not here model any of these possible sources of spurious variations in the Solow residual.

Formalization of this idea requires that we introduce a new variable, namely the “long run” level of technology expected at t , v_t . In direct analogy with (1), we define it as

$$\log v_t \equiv \lim_{T \rightarrow \infty} E_t[\log z_{t+T} - T \log \gamma_z] \quad (25)$$

where z_t continues to denote the technology factor in the period t aggregate production function, and γ_z is the unconditional average rate of growth of z . The variable v_t can be thought of as the level of basic knowledge about technological opportunities that the society has at t . Actual technology can differ from this because of the time taken for new technologies to be adopted. To be consistent with (25), v_t must follow a random walk. The reason is that the revisions in the expectation of the level of expected long run technology must be independent of any information available at t . Thus, we have that

$$\log v_t = \log v_{t-1} + \epsilon_t \quad (26)$$

Our analysis in section 3 assumed that z_t was always equal to v_t . Here we propose to relax this assumption, but in the simplest possible way. Thus we assume that

$$\log z_t = \beta \log z_{t-1} + (1 - \beta) \log v_t \quad (27)$$

This equation implies that

$$\hat{\gamma}_t^z = \beta \hat{\gamma}_{t-1}^z + (1 - \beta) \epsilon_t \quad (28)$$

so that the rate of growth of technology follows a first-order AR process. It follows that all conditional expectations of the form $E_t[\log z_{t+j}]$ are functions of only two state variables, z_t and v_t .

This equation defines a one-parameter family of stochastic processes for technology, with the case considered in section 3 corresponding to $\beta = 0$. We now consider values ranging over the interval $0 \leq \beta < 1$, where higher β corresponds to slower diffusion of the innovation. Mansfield (1968) contains estimates of rates of diffusion for many innovations.⁴⁰ He shows that the time elapsed before half the firms in an industry adopt a major innovation has varied between 1 and 15 years. We thus consider a rate of diffusion such that half

⁴⁰Mansfield (1968) stresses that the stock of adopters follows an S-curve and one could view this as evidence that the stochastic process for z is more complex than in (27); an S-shaped impulse response of z to an innovation requires that the growth rate of z be an AR of at least second order. We do not pursue this further here, for two reasons. The first is the scarcity of information on the basis of which to calibrate several different parameters. The second is that the S-curve discussed in the empirical literature is an *ex post* description of successful innovations. Because unsuccessful innovations probably start out looking similar to successful ones, it is hard to know at which point in the S-curve one should imagine that people know that the long-run technology will be different. This idea fits well with Ellison and Fudenberg's (1993) explanation of slow diffusion based on slow learning.

the change in v is embedded in a change in z after 30 quarters. This implies that β is equal to about 0.98. We thus give special emphasis to this value.

In this variant of the model, output, consumption and hours depend on four variables, the current capital stock, N_t , the current level of technology z_t and the eventual level of technology v_t . But, once again it is straightforward to show that consumption and output are homogeneous of degree one in K_t , $z_t N_t$ and $v_t N_t$, while h_t is homogeneous of degree zero in this same set of variables. This means that the transformed variables $\frac{Y_t}{z_t N_t}$, $\frac{C_t}{z_t N_t}$ and h_t depend on two variables, namely κ_t and the ratio of v_t to z_t (or equivalently, $(\log v_t - \log z_t)$). Since C_t/Y_t and h_t depend just on these two state variables, it is generally possible to reconstruct the state variables from the two stationary variables included in our VAR. Thus the theoretical model implies that the forecastable changes estimated by our VAR should correspond to the forecastable changes given the information set of agents in the model.⁴¹

Table 12 presents the theoretical predictions for a model with our baseline preference specifications, but for alternative values of β . These predictions are computed keeping the standard deviation of ϵ_t , σ_ϵ , equal to 0.00732. The table shows that the R^2 , the fraction of changes in output over 12 quarters that are predictable, rises with β . This is due to the fact that higher values of β make the future rate of technical progress more predictable. While this increase in R^2 seems desirable, increases in β also lower the standard deviation of predictable output changes. The reason is that higher values of β make the rate of growth of technology smoother (since it becomes more serially correlated as β is increased, in the limit approaching a random walk). Thus, for a given σ_ϵ , the variance of changes in $\log z$ falls, and consequently the variance of changes in $\log Y$ as well. This reduction in the variability of output also leads to a reduction in the amount of predictable variability. Thus slow diffusion of technology does not lead to larger forecastable movements in output.

Table 12 also shows that many of the correlations between output growth and other variables that were problematic in the case of $\beta = 0$ remain problematic for higher values of β . In particular, high expected growth is still correlated with high values of hours and low values of C/Y . This is surprising, because consumption rises more than output immediately following a positive innovation if β is large. This can be

⁴¹ Thus the possibility of mis-identification of technology shocks from the VAR residuals that Lippi and Reichlin show can arise in the case of slow diffusion of technical progress does not occur in this model. Hence for purposes of calibrating this model, we are again able to take the estimated variance of innovations in the long-run forecast of output from the VAR as the variance of innovations in v_t .

seen in Figure 6, which shows the impulse responses of consumption, output and hours to a unit innovation ϵ_t when β is equal to 0.98. The reason consumption jumps so much is that there continues to be a strong wealth effect of the innovation, even though the productivity of existing inputs has increased very little. But, while C/Y rises immediately, it soon falls, and spends most of the transition period below its steady state value. For that reason, the overall correlation between output growth and C/Y continues to be negative.

With this high value of β , the regression coefficients of expected consumption growth and expected hours growth on expected output growth are closer to their empirical counterparts. The first is now below 1, as in the data, while the second is now larger. On the other hand, the model still predicts that labor productivity should rise together with output. This prediction seems hard to avoid in models where the transitory dynamics are due to persistent technology shocks that cause long run growth. It is probably the single biggest reason for our feeling that the shocks that lead to long run growth do not seem capable of generating the sort of predictable output movements that we have been exploring in this paper.

Thus allowing for forecastable technical progress of this particular kind does not help to explain the size or character of the forecastable variations in output growth. What is more, in the case of a high value of β , the model implies that technology shocks account for only a trivial fraction of the *overall* variability of output growth as well. Whereas in the case $\beta = 0$, the predicted standard deviation of output growth is .0053 (nearly half the standard deviation of actually observed output growth), in the case $\beta = .98$, the predicted standard deviation of output growth is only .0009. This allows us to sharpen the point made earlier: Accepting the existence of a stochastic trend does not require us to believe that innovations in that trend play any significant role in the generation of business cycles. It is now clear that this is true not only when by “business cycles” we mean forecastable output movements; innovations in the trend need not play any significant role in the generation of period to period variability in aggregate quantities, forecastable or otherwise. If the correct model were of this kind, but with additional independent disturbances in addition to the technology shocks, essentially all of the variability in aggregate quantities would have to be due to the other shocks. Thus the existence of a “unit root” in output does not, in itself, imply anything about the role of technology shocks in generating output variability.⁴²

⁴²The variance decompositions reported by K-P-S-W do, of course, provide further evidence in this regard. The amount of the variance in output and other quantities over, say, a 12-quarter horizon that they attribute to the innovations in the “balanced-growth shock” is certainly not consistent with a theoretical model like the one discussed here, with a high value of β but with additional transitory disturbances.

8 Conclusions

We have demonstrated that the forecastable movements in output – what we would argue is the essence of the “business cycle” – are inconsistent with a standard growth model disturbed solely by random shocks to the rate of technical progress. In the case of a standard calibration of parameter values, the model predicts neither the magnitude of these forecastable changes nor their basic features, such as the signs of the correlations among the forecastable changes in various aggregate quantities. We have also argued that contemplation of parameter values outside the range typically assumed in the real business cycle literature does little to improve the model’s performance in this regard, while significantly worsening the model’s performance on dimensions emphasized in that literature.

Various possible interpretations might be given for the failure of this particular type of stochastic growth model to explain the business cycle. It may be that the business cycle is mainly caused by disturbances other than technology shocks, that the model errs in its account of the dynamic response to technology shocks, or that the technology shocks that account for the business cycle have serial correlation properties very different from those assumed here. Whichever possibility turns out to account for more of the failure, one can conclude that the standard growth model provides a poor description of the “propagation mechanism” by which the effects of shocks evolve over time. For we can show that the mere introduction of additional disturbances to the equilibrium conditions of the model cannot solve the problem, regardless of the nature or magnitude of the disturbances contemplated, if these additional disturbances are purely transitory. Thus the additional disturbances (whether they represent additional transitory components of the productivity factor, or shocks of some other kind) would have to exhibit significant persistence, and the mechanism by which these disturbances persist over many quarters would turn out to be a crucial source of business cycle dynamics – in essence, a propagation mechanism in addition to those present in the basic growth model.

But it is not obvious that one should assume that the equations of the basic model are correct except for the absence of stochastic disturbance terms. Quite possibly, the standard growth model must be modified to include other sources of dynamics before it can be used to model business cycles. Some obvious candidates would include inventory dynamics, slow adjustment of the work force as is implied in models of “labor hoarding,” or slow adjustment of nominal wages and/or prices as is implied by models with overlapping contracts or costs of price adjustment. The degree to which mechanisms of these sorts might better account

for the size and nature of the forecastable movements in output and other variables documented here remains a topic for future research.

9 References

- Beveridge, Stephen and Charles R. Nelson: "A New Approach to the Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the 'Business Cycle'," *Journal of Monetary Economics* 7, March 1981, 151-74.
- Blanchard, Olivier J. and Danny Quah: "The Dynamic Effects of Aggregate Supply and Demand Disturbances," *American Economic Review*, 79, September 1989, 655-73.
- Blanchard, Olivier J. and Stanley Fischer: *Lectures on Macroeconomics*, Cambridge: MIT Press, 1989.
- Brock, William A. and Leonard J. Mirman, "Optimal Economic Growth and Uncertainty: The Discounted Case," *Journal of Economic Theory*, 4, 1972, 497-513.
- Campbell, John Y.: "Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis," *Econometrica*, 55, November 1987, 1249-73.
- and N. Gregory Mankiw: "Consumption, Income and Interest Rates: Reinterpreting the Time series Evidence," *NBER Macroeconomics Annual 1989*, 185-216.
- Christiano, Lawrence J. and Martin Eichenbaum: "Current Real-Business-Cycle Theories and Aggregate Labor Market Fluctuations," *American Economic Review*, 82, June 1992, 430-50.
- Cochrane, John H.: "Permanent and Transitory Components of GNP and Stock Prices," *Quarterly Journal of Economics*, 1994a, forthcoming.
- Cochrane, John H.: "Shocks," *Carnegie-Rochester Conference*, 1994b, forthcoming.
- and Argia Sbordone: "Multivariate Estimates of the Permanent Components of GNP and Stock Prices," *Journal of Economic Dynamics and Control*, 12, June/September 1988, 255-96.
- Cogley, Timothy and James M. Nason: "Do Real Business Cycle Pass the Nelson-Plosser Test?," mimeo, August 1992.
- Cooley, Thomas F. and Gary D. Hansen, "Money and the Business Cycle," mimeo, October 1993. Forthcoming in T.F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press.
- Ellison, Glenn and Drew Fudenberg: "Rules of Thumb for Social Learning," *Journal of Political Economy*,

101, August 1993, 612-43.

Evans, George: "Output and Unemployment Dynamics in the United States: 1950-1985," *Journal of Applied Econometrics*, 4, July-September 1989, 213-37.

Evans, George and Lucrezia Reichlin: "Information, Forecasts and Measurement of the Business Cycle," mimeo, OFCE, Paris, May 1993.

Friedman, Benjamin and Kenneth N. Kuttner: "Money, Income, Prices and Interest Rates," *American Economic Review*, 82, 472-92.

Gordon, Robert J.: "The 'End-of-Expansion' Phenomenon in Short Run Productivity Behavior," *Brookings papers on Economic Activity*, (1979), 447-63.

Jovanovic, Boyan, and Saul Lach, "Diffusion Lags and Aggregate Fluctuations," *NBER Working Paper No. 4455*, September 1993.

King, Robert G., Charles I. Plosser and Sergio Rebelo: "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," *Journal of Monetary Economics*, 21, March 1988a, 195-232.

———"Production, Growth and Business Cycles: II. New Directions," *Journal of Monetary Economics*, 21, March 1988b, 309-342.

King, Robert G., Charles I. Plosser, James H. Stock and Mark W. Watson: "Stochastic Trends and Economic Fluctuations," *American Economic Review*, 81, September 1991, 819-40.

King, Robert G. and Sergio Rebelo: "Low Frequency Filtering and Real Business Cycles," *Journal of Economic Dynamics and Control*, 17, 1993, 207-31.

Kydland, Finn E. and Edward C. Prescott: "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, (November 1982): 1345-70.

Lippi, Marco and Lucrezia Reichlin: "The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment," *American Economic Review*, 83, June 1993, 644-52.

Mansfield, Edwin: *The Economics of Technological Change*, New York: Norton, 1968.

Nelson, Charles R. and Charles I. Plosser: "Trends and Random Walks in Macroeconomic Time Series: Some Evidence and Implications," *Journal of Monetary Economics*, 10, 1982, 139-62.

Plosser, Charles I.: "Understanding Real Business Cycles," *Journal of Economic Perspectives*, 3, Summer 1989, 51-78.

- Prescott, Edward C.: "Theory Ahead of Business Cycle Measurement," *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall 1986, 9-22.
- Rotemberg, Julio J. and Michael Woodford: "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100, December 1992, 1153-207.
- and —: "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets," forthcoming in T. Cooley, ed. *Frontiers of Business Cycle Research*, Princeton: Princeton University Press, 1994
- Shapiro, Matthew D. and Mark W. Watson: "Sources of Business Cycle Fluctuations," *NBER Macroeconomics Annual*, 1988, 111-148.
- Sims, Christopher A.: "Output and Labor Input in Manufacturing," *Brookings Papers on Economic Activity*, III, (1974), 695-737.
- Watson, Mark W.: "Measures of Fit for Calibrated Models," *Journal of Political Economy*, 101, December 1993, 1011-41.

Table 1 Regression Results

Explanatory vars.	Δy	$(c - y)$	h	$\Delta^2 y$	$\Delta(c - y)$	Δh
Constant	0.023 0.016	-0.042 0.014	0.010 0.012	0.005 0.001	-0.044 0.013	0.344 0.080
Δy_{-1}	0.570 0.167	-0.469 0.147	0.570 0.127	-0.696 0.098		
Δy_{-2}	0.002 0.087	0.017 0.077	-0.005 0.066			
$(c - y)_{-1}$	0.663 0.166	0.330 0.146	0.490 0.126		-0.098 0.030	
$(c - y)_{-1}$	-0.618 0.158	0.568 0.139	-0.458 0.120			
h_{-1}	0.215 0.128	-0.283 0.113	1.450 0.097			-0.079 0.019
h_{-2}	-0.314 0.135	0.316 0.119	-0.503 0.102			
$\Delta^2 y_{-1}$				0.013 0.091		
$\Delta^2 y_{-2}$				0.129 0.075		
$\Delta(c - y)_{-1}$					0.122 0.074	
$\Delta(c - y)_{-2}$					0.126 0.075	
Δh_{-1}						0.760 0.070
Δh_{-2}						-0.135 0.074
Trend						3e-4 8e-5

Data from 1948.4 to 1993.2. Standard Errors below coefficient estimates.

Table 2
Unconditional Second Moments

Moment	Baseline Model	U.S. Data
$S.D.\Delta y$	0.0053	0.0113
$S.D.\Delta c/S.D.\Delta y$	0.585	0.478
$S.D.\Delta i/S.D.\Delta y$	2.019	2.911
$S.D.\Delta h/S.D.\Delta y$	0.358	0.934
$S.D.\Delta p/S.D.\Delta y$	0.660	0.690
$Corr.(\Delta c, \Delta y)$	0.987	0.5223
$Corr.(\Delta h, \Delta y)$	0.977	0.7475
$Corr.(\Delta i, \Delta y)$	0.994	0.9452
$Corr.(\Delta p, \Delta y)$	0.994	0.4373

Table 3
Estimated standard deviations of forecasted changes

Horizon	$\widehat{\Delta y}^k$	$\widehat{\Delta c}^k$	$\widehat{\Delta i}^k$	$\widehat{\Delta h}^k$	$\widehat{\Delta p}^k$	$\Delta y^k - \widehat{\Delta y}^k$	R^2
1.	0.0062 0.0007	0.0023 0.0004	0.0175 0.0021	0.0077 0.0005	0.0037 0.0005	0.0094	0.308
2.	0.0108 0.0014	0.0044 0.0007	0.0288 0.0039	0.0130 0.0013	0.0059 0.0009	0.0148	0.347
4.	0.0190 0.0029	0.0082 0.0014	0.0494 0.0077	0.0204 0.0029	0.0081 0.0014	0.0220	0.429
8.	0.0299 0.0039	0.0139 0.0022	0.0767 0.0096	0.0309 0.0045	0.0080 0.0017	0.0273	0.550
12	0.0326 0.0034	0.0172 0.0027	0.0825 0.0068	0.0341 0.0031	0.0072 0.0018	0.0304	0.537
24	0.0309 0.0037	0.0224 0.0046	0.0772 0.0053	0.0323 0.0012	0.0067 0.0020	0.0394	0.384
∞	0.0310 0.0044	0.0272 0.0089	0.0585 0.0039	0.0325 0.0000	0.0068 0.0023		0.000

Standard Errors based on coefficient uncertainty below estimates.

Table 4
Estimated correlations of predicted output changes

with the regressors			
Regressor:	Δy_t	$(c_t - y_t)$	h_t
Horizon:			
1.	0.597	0.214	-0.599
	0.090	0.109	0.090
2.	0.530	0.312	-0.793
	0.095	0.132	0.077
4.	0.291	0.451	-0.960
	0.095	0.142	0.030
8.	0.009	0.533	-0.995
	0.062	0.127	0.007
12.	-0.087	0.546	-0.979
	0.058	0.119	0.013
24.	-0.093	0.544	-0.978
	0.058	0.171	0.012
∞	-0.092	0.541	-0.978
	0.071	0.240	0.014

Standard Errors based on coefficient uncertainty below estimates.

Table 5
Estimated regression coefficients among forecasted changes

	$\widehat{\Delta c}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta i}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta h}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta p}_t^k$ on $\widehat{\Delta y}_t^k$
Horizon:				
1.	0.2537 0.0590	2.7412 0.1376	1.0922 0.1102	-0.0922 0.1102
2.	0.3119 0.0603	2.6057 0.1406	1.0741 0.1154	-0.0741 0.1154
4.	0.3443 0.0627	2.5299 0.1463	0.9881 0.1079	0.0119 0.1079
8.	0.3678 0.0595	2.4751 0.1389	0.9964 0.1098	0.0036 0.1098
12.	0.3997 0.0614	2.4008 0.1434	1.0230 0.1134	-0.0230 0.1134
24.	0.4970 0.1212	2.1738 0.2828	1.0253 0.1215	-0.0253 0.1215
∞	0.5711 0.2358	2.0008 0.5502	1.0252 0.1375	-0.0252 0.1375

Standard Errors based on coefficient uncertainty below estimates.

Table 6
The Calibrated Parameters

Parameter	Defined by	Values	Description
γ_z		1.004	Steady state growth rate of technology (per quarter)
γ_n		1.004	Labor force growth rate (per quarter)
s_C	$\frac{\bar{C}}{\bar{Y}}$	0.70	Share of private consumption expenditure in private value added net of govt. purchases
δ		0.025	Rate of depreciation of capital stock (per quarter)
s_K	$\frac{F_K \bar{K}}{\bar{F}}$	0.42	Share of capital costs in total costs
ϵ_{KL}	$\frac{F_K F_L}{F_{KL} \bar{F}}$	1	Elasticity of substitution between capital and hours
r	$F_K - \delta$ or $\gamma_z^\sigma \beta^{-1} - 1$	0.01625	Steady state real rate of return (per quarter)
$1/\sigma$		1	Intertemporal elasticity of substitution of consumption holding hours worked constant
ϵ_{HW}		4	Intertemporal elasticity of labor supply
σ_ϵ		.00732	Standard deviation of permanent technology shock

Note: Except for rate of population growth, parameters displayed above μ take the same values as in King, Plosser and Rebelo (1988a).

Table 7
Predicted standard deviations of forecasted changes

Horizon:	$\widehat{\Delta y}^k$	$\widehat{\Delta c}^k$	$\widehat{\Delta i}^k$	$\widehat{\Delta h}^k$	$\widehat{\Delta p}^k$	$\Delta y^k - \widehat{\Delta y}^k$	R^2
1	0.0003	0.0007	0.0005	0.0003	0.0006	0.0053	0.0035
2	0.0006	0.0013	0.0010	0.0006	0.0012	0.0076	0.0065
4	0.0012	0.0025	0.0019	0.0011	0.0022	0.0109	0.0113
8	0.0021	0.0046	0.0036	0.0020	0.0041	0.0159	0.0176
12	0.0029	0.0063	0.0049	0.0027	0.0056	0.0200	0.0208
24	0.0045	0.0097	0.0076	0.0042	0.0087	0.0300	0.0224
∞	0.0066	0.0141	0.0110	0.0060	0.0126		0.0000

Table 8
Predicted regression coefficients among forecasted changes

$\widehat{\Delta c}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta i}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta h}_t^k$ on $\widehat{\Delta y}_t^k$	$\widehat{\Delta p}_t^k$ on $\widehat{\Delta y}_t^k$
2.1464	-1.6750	-0.9198	1.9198

Table 9
The Effect of Varying the Preference Parameters

σ	ϵ_{HW}	S.D. $\widehat{\Delta y}^{12}$	Corr ($\widehat{\Delta y}^{12}, c - y$)	Corr($\widehat{\Delta y}^{12}, h$)	$\widehat{\Delta c}_t^{12}$ on $\widehat{\Delta y}_t^{12}$	$\widehat{\Delta h}_t^{12}$ on $\widehat{\Delta y}_t^{12}$	$\widehat{\Delta p}_t^{12}$ on $\widehat{\Delta y}_t^{12}$
0.6	0.2	0.0041	-1.0	1.0	2.0990	-0.1799	1.1799
0.6	1.0	0.0030	-1.0	1.0	3.1763	-1.0312	2.0312
0.6	4.0	0.0003	-1.0	1.0	39.0290	-27.9527	28.9527
0.6	10.0	0.0025	1.0	-1.0	-5.3574	5.2520	-4.2520
1	0.2	0.0039	-1.0	1.0	1.5827	-0.0978	1.0978
1	1.0	0.0035	-1.0	1.0	1.7813	-0.3918	1.3918
1	4.0	0.0029	-1.0	1.0	2.1460	-0.9193	1.9193
1	10.0	0.0026	-1.0	1.0	2.3765	-1.2504	2.2504
1	∞	0.0023	-1.0	1.0	2.6845	-1.6862	2.6862
4	0.2	0.0031	1.0	-1.0	0.8975	0.0621	0.9379
4	1	0.0029	1.0	1.0	0.8854	-0.0488	1.0488
4	4	0.0029	1.0	1.0	0.8867	-0.0365	1.0365
4	10	0.0030	1.0	1.0	0.8869	-0.0348	1.0348
4	∞	0.0030	1.0	1.0	0.8870	-0.0337	1.0337
10	0.2	0.0022	1.0	1.0	0.6531	-0.0425	1.0425
10	1	0.0022	1.0	1.0	0.6580	-0.0281	1.0281
10	4	0.0022	1.0	1.0	0.6586	-0.0264	1.0264
10	10	0.0022	1.0	1.0	0.6587	-0.0261	1.0261
10	∞	0.0022	1.0	1.0	0.6588	-0.0259	1.0259

Table 10
Forecasts for k Periods Starting in Period $t + j$

j	k	S.D. $\widehat{\Delta y}_j^k$	Corr($\widehat{\Delta y}_j^k, (c - y)_j$)	Corr($\widehat{\Delta y}_j^k, h_j$)	$\widehat{\Delta c}_j^k$ on $\widehat{\Delta y}_j^k$	$\widehat{\Delta h}_j^k$ on $\widehat{\Delta y}_j^k$	$\widehat{\Delta p}_j^k$ on $\widehat{\Delta y}_j^k$
1	1	0.0051	0.284	-0.838	0.349	1.005	-0.005
		0.0008	0.160	0.087	0.063	0.116	0.116
	2	0.0099	0.372	-0.920	0.349	0.980	0.020
		0.0017	0.158	0.053	0.062	0.109	0.109
	4	0.0186	0.469	-0.987	0.349	0.991	0.009
		0.0030	0.148	0.015	0.059	0.108	0.108
	8	0.0286	0.533	-0.998	0.364	1.050	-0.050
		0.0035	0.127	0.004	0.056	0.114	0.114
	12	0.0303	0.548	-0.989	0.396	1.080	-0.080
		0.0033	0.125	0.009	0.062	0.118	0.118
	24	0.0286	0.546	-0.988	0.492	1.082	-0.082
		0.0037	0.189	0.009	0.133	0.127	0.127
	∞	0.0288	0.543	-0.988	0.564	1.082	-0.082
		0.0043	0.261	0.010	0.256	0.145	0.145
2	8	0.0257	0.501	-0.999	0.362	1.111	-0.111
		0.0033	0.135	0.004	0.057	0.122	0.122
	∞	0.0252	0.510	-0.993	0.566	1.142	-0.142
		0.0042	0.312	0.011	0.314	0.1569	0.157
4	8	0.0178	0.348	-0.999	0.376	1.223	-0.223
		0.0034	0.183	0.005	0.079	0.146	0.146
	∞	0.0166	0.355	-0.995	0.619	1.251	-0.251
		0.0043	0.536	0.015	0.632	0.202	0.202
8	8	0.0046	0.066	-0.995	0.407	1.455	-0.455
		0.0028	0.791	0.046	0.979	0.188	0.188
	∞	0.0038	0.089	-0.974	0.680	1.476	-0.476
		0.0025	2.463	0.168	8.921	0.189	0.189
12	8	0.0016	0.517	-0.988	-0.955	0.977	0.023
		0.0025	1.720	0.143	4.617	0.554	0.554
	∞	0.0016	0.453	-0.957	-2.191	0.988	0.012
		0.0032	4.250	0.142	25.65	1.041	1.041
24	8	0.0001	-0.718	-0.965	13.11	0.998	0.002
		0.0011	5.846	1.159	37.90	5.316	5.316
	∞	0.0002	-0.871	-0.983	26.17	1.026	-0.026
		0.0031	4.826	0.477	330.9	10.57	10.57

Standard Errors based on coefficient uncertainty below estimates.

Table 11
Predicted Standard Deviation of k Period Output Changes Starting at $t + j$

$j =$ $k:$	1	2	4	8	12	24
8	0.0020	0.0019	0.0018	0.0014	0.0012	0.0007
∞	0.0063	0.0060	0.0054	0.0045	0.0037	0.0020

Table 12
The Effect of Slow Diffusion of Technical Progress

β	S.D. $\widehat{\Delta Y}^{12}$	R^2	$\text{Corr}(\widehat{\Delta y}^{12}, c - y)$	$\text{Corr}(\widehat{\Delta y}^{12}, h)$	$\widehat{\Delta c}_t^{12}$ on $\widehat{\Delta y}_t^{12}$	$\widehat{\Delta h}_t^{12}$ on $\widehat{\Delta y}_t^{12}$	$\widehat{\Delta p}_t^{12}$ on $\widehat{\Delta y}_t^{12}$
0.99	0.0001	0.990	-0.535	0.523	0.885	0.091	0.908
0.98	0.0002	0.975	-0.394	0.393	0.806	0.155	0.845
0.95	0.0006	0.887	-0.213	0.213	0.685	0.252	0.748
0.9	0.0012	0.695	-0.125	0.125	0.630	0.296	0.704
0.8	0.0022	0.409	-0.132	0.132	0.634	0.293	0.707
0.7	0.0027	0.256	-0.205	0.205	0.676	0.259	0.741
0.6	0.0030	0.170	-0.304	0.304	0.746	0.203	0.797
0.5	0.0032	0.118	-0.420	0.420	0.849	0.121	0.879
0.4	0.0033	0.083	-0.550	0.550	0.992	0.006	0.994
0.3	0.0033	0.059	-0.688	0.688	1.191	-0.153	1.153
0.2	0.0032	0.041	-0.826	0.826	1.459	-0.367	1.367
0.1	0.0030	0.029	-0.945	0.945	1.798	-0.638	1.638

Figure 1
Expected Declines in Private Output and NBER Troughs




Figure 2
Expected Declines in Next 12 Quarters and Detrended Output




Figure 3
A 1% Shock to z ; $\sigma = 1$; $\epsilon_{HW} = 4$




Figure 4
 A 1% Shock to z ; $\sigma = 0.6$; $\epsilon_{HW} = 10$




Figure 5
A 1% Shock to z ; $\sigma = 4$; $\epsilon_{HW} = 0.2$





Figure 6
Impulse Response Function with Slow Diffusion: $\beta = 0.98$



MIT LIBRARIES



3 9080 00856629 8

2267 027

Date Due

Lib-26-67

